

Metrics for Multivariate Dictionaries

Sylvain Chevallier

LISV - Université de Versailles Saint-Quentin
78140 Velizy, France

Quentin Barthélemy

CEA, LIST - LOAD
91191 Gif-sur-Yvette, France

Jamal Atif

LRI - Université Paris-Sud 11
91410 Orsay, France

`sylvain.chevallier@uvsq.fr`, `quentin.barthelemy@cea.fr`,
`jamal.atif@lri.fr`

February 27, 2013

Abstract

Overcomplete representations and dictionary learning algorithms kept attracting a growing interest in the machine learning community. This paper addresses the emerging problem of comparing multivariate overcomplete representations. Despite a recurrent need to rely on a distance for learning or assessing multivariate overcomplete representations, no metrics in their underlying spaces have yet been proposed. Henceforth we propose to study overcomplete representations from the perspective of frame theory and matrix manifolds. We consider distances between multivariate dictionaries as distances between their spans which reveal to be elements of a Grassmannian manifold. We introduce Wasserstein-like set-metrics defined on Grassmannian spaces and study their properties both theoretically and numerically. Indeed a deep experimental study based on tailored synthetic datasets and real EEG signals for Brain-Computer Interfaces (BCI) have been conducted. In particular, the introduced metrics have been embedded in clustering algorithm and applied to BCI Competition IV-2a for dataset quality assessment. Besides, a principled connection is made between three close but still disjoint research fields, namely, Grassmannian packing, dictionary learning and compressed sensing.

Keywords: Dictionary Learning, Metrics, Frames, Grassmannian Manifolds, Packing, Compressed Sensing, Multivariate Dataset.

1 Introduction

A classical question in machine learning is how to choose a good feature space to analyze the input data. One possible and elegant response is to learn this feature space over mild hypotheses. This approach is known as dictionary learning and follows the dictionary-based framework which have provided several important results in the last decades (Meyer, 1995; Mallat, 1999; Tropp, 2004; Candès et al., 2006; Gribonval and Schnass, 2010; Jenatton et al., 2012). In dictionary-based approaches, an expert selects a specific family of basis functions, called atoms, known to capture important features of the input data, for example wavelets (Mallat, 1999) or curvelets (Candès and Donoho, 2004). When no specific expert knowledge is available, dictionary learning algorithms could learn those atoms from a given dataset. The problem is then stated as an optimization procedure usually under some sparsity constraints. Thanks to the pioneering works on sparsity constraint decomposition of Mallat and Zhang (1993), on the Lasso problem (Tibshirani, 1996) and on sparse signal recovery (Candès et al., 2006), efficient algorithms are available both for projection on overcomplete representations (Efron et al., 2004; Beck and Teboulle, 2009) and dictionary learning (Engan et al., 2000; Aharon et al., 2006; Mairal et al., 2010).

Despite the very profusion of research papers dealing with dictionary-based approaches and overcomplete representations in general, except some noticeable exceptions (Balan, 1999; Skretting and Engan, 2010), few results have been reported on how the constructed representations should be compared, and in a more theoretical standpoint, on which topological space these representations live.

Thus, to qualitatively assess a specific dictionary learning algorithm, one has to indirectly evaluate it through a benchmark based on a task performance (Engan et al., 2000; Grosse et al., 2007; Mairal et al., 2009; Skretting and Engan, 2010). Meanwhile, one can find in the literature some hints for dictionaries comparison with the aim of learning assessment (Aldroubi, 1995; Aharon et al., 2006; Vainsencher et al., 2011; Skretting and Engan, 2011), but they fall short to define a true metric. A recurrent view to dictionary comparison is based on the cross-gramian matrix and is only defined on univariate dictionaries. In Lesage (2007), it is called a correlogram and measures, such as precision and recall, were defined. In Aldroubi (1995), dictionaries are compared through the analysis Gram operator induced norm. The work of Balan (1999) introduces a mapping from the frame sets to a continuous functional space. The constructed functions allow then for the definition of an equivalence class and a partial order. Nonetheless this approach is not invariant of linear transformation between frame elements.

In harmonic analysis, overcomplete dictionaries are called frames (Christensen, 2003). Frames are tightly linked with the work of Gabor, as they were formally defined by Duffin and Schaeffer (1952) and later popularized

by Daubechies et al. (1986) to be widely studied as a theoretical framework for wavelets (Mallat, 1999). They have been used in numerous domains, for example signal processing, compression or sampling. They have also been actively investigated for solving packing problems (Conway et al., 1996; Strohmer and Heath Jr., 2003; Dhillon et al., 2008). In its original formulation by Tammes (1930), the packing problem in a compact metric space consists in finding the best way to embed points on the surface of a sphere such that each pair of point is separated by the largest distance as possible. When considering the packing of subspaces, which has applications in wireless communications, it has been the topic of a very active community, see for example (Casazza and Kutyniok, 2004; Kutyniok et al., 2009; Strawn, 2012).

Of particular interest is the remarkable equivalence between the Equiangular Tight Frames design and the Grassmannian packing problem (Strohmer and Heath Jr., 2003). This opens new exciting perspectives on exploiting the theoretical results on Grassmannian manifolds and bring them to the context of frame theory and inherently overcomplete dictionaries.

Several metrics have been proposed in Grassmannian spaces, such as the chordal distance (Conway et al., 1996; Hamm and Lee, 2008), the projection-norm (Edelman et al., 1999) or the Binet-Cauchy distance (Wolf and Shashua, 2003; Vishwanathan and Smola, 2004). These distances have been used to define a reproducing kernel Hilbert space, for SVM-based approaches or to find the best packing with a given subspace, and were also used in image processing, see Lui (2012) for a review.

The Grassmannian spaces define a natural framework to work with multicomponent dataset. This opens new approaches to account for finding spatial relations in multivariate time-series or in multichannel images. Previous works on dictionary learning in a multivariate context are applied on multimodal audio-visual data (Monaci et al., 2007, 2009), electro-cardiogram signals (Mailhé et al., 2009), color images (Mairal et al., 2008), hyperspectral images (Moudden et al., 2009), stereo images (Tošić and Frossard, 2011b), 2D spatial trajectories (Barthélemy et al., 2012) and electro-encephalogram (EEG) signals (Barthélemy et al., 2013a).

Main contribution of the paper is to introduce metrics suited to multivariate dictionaries, which are invariant to linear combinations (see Section 5). To define these metrics, intermediate results were necessary. In particular, transport metrics are defined over Grassmannian manifolds. A principled connection has been made between multivariate dictionaries, frame theory and Grassmannian manifolds by exploiting the connection between the frame design problem and the Grassmannian packing (Section 4).

In the sequel, we first begin by recalling some definitions about Grassmannian manifolds and their metrics in Section 2, supplementary material on matrix manifolds are given in Appendix A. In Section 3, necessary material from frame theory is recalled. An original formalized connection be-

tween Grassmannian frames, compressed sensing and dictionary learning is detailed in Section 4. In Section 5, distances between overcomplete representations are constructed as Wasserstein-like set metrics where ground metrics are those defined on Grassmannian manifolds. In Section 6, dictionary learning principle is detailed for multivariate data, with different models. Assessment on synthetic data is presented in Section 7 with experimental results for several multivariate dictionary learning algorithms, that is overcomplete representations on multichannel dataset. Proposed metrics allow to estimate the empirical convergence of dictionary learning algorithms and thus to evaluate the contribution of these different algorithms. Experiences on real dataset are shown in Section 8, exploiting the overcomplete representations to assess the quality of training dataset through unsupervised analyses. A well known multivariate dataset is investigated, namely the Brain-Computer Interface Competition IV-2a. An evaluation of the subject variability and of the difference between the training and testing datasets is produced, hence offering novel insights on the dataset consistency. Section 9 concludes this paper and points out some future research directions.

2 Grassmannian Manifolds and their Metrics

This section recalls some definitions from algebraic geometry that will be helpful to characterize the metrics associated with Grassmannian manifolds. Some basic definitions, including the Grassmannian manifold one, is given in Appendix A.

2.1 Preliminaries

We consider here an n -dimensional vector space \mathcal{V} with inner product $\langle \cdot, \cdot \rangle$ defined over the field \mathbb{R} . We will denote by u, w the elements (vectors) of \mathcal{V} and by U, W the matrices of dimension $n \times \rho$ defined as $U = [u_1, \dots, u_\rho]$, $W = [w_1, \dots, w_\rho]$. The transpose of a vector u (or a matrix U) is denoted u^T (U^T). The inner product induces the ℓ_2 norm $\|u\|_2^2 = \langle u, u \rangle$. The Frobenius norm is defined as $\|U\|_F^2 = \text{trace}(U^T U)$ and its associated inner product is defined as $\langle U, W \rangle_F = \text{trace}(W^T U)$. The subspace spanned by U is denoted as \mathcal{U} and its dimension (rank of U) as ϱ , thus $\dim(\mathcal{U}) = \text{rank}(U) = \varrho$.

2.2 Principal Angles

A central notion to characterize the distance between two subspaces is the *principal angles*, which is easy to interpret as shown on Figure 1.

Definition 1 *The principal angles $0 \leq \theta_1 \leq \dots \leq \theta_\varrho \leq \frac{\pi}{2}$ between two*




Figure 1: Illustration of the principal angles between two subspaces \mathcal{U} and \mathcal{W} . In this example, $\rho = \varrho = 2$, θ_1 is null and θ_2 is plotted in the figure.

subspaces \mathcal{U} and \mathcal{W} are defined recursively by

$$\begin{aligned} \cos \theta_k &= \max_{y_k \in \mathcal{U}} \max_{z_k \in \mathcal{W}} y_k^T z_k \\ \text{s.t. } & \|y_k\|_2 = 1, \|z_k\|_2 = 1, \\ \text{and } & y_k^T y_j = 0, z_k^T z_j = 0, j = 0, \dots, k-1. \end{aligned}$$

It should be noticed that principal angles are also defined when the dimension of the two subspaces are different, for example $r = \dim \mathcal{W} \geq \dim \mathcal{U} = \varrho \geq 1$, in which case the maximum number of principal angles is equal to the dimension of the smallest subspace.

A more computationally tractable definition of the principal angles relies on the singular value decomposition of two bases $A = \{a_i\}_{i \in I}$ spanning \mathcal{U} and $B = \{b_i\}_{i \in J}$ spanning \mathcal{W} , with I and J two indexing sets:

$$A^T B = \underline{Y} \Sigma \underline{Z}^T = \underline{Y} (\cos \boldsymbol{\theta}) \underline{Z}^T,$$

where $\underline{Y} = [y_1, \dots, y_\varrho]$, $\underline{Z} = [z_1, \dots, z_\varrho]$. We denote by $\boldsymbol{\theta}$ the ϱ -vector formed by the principal angles θ_k , $k = 1, \dots, \varrho$. Here $\cos \boldsymbol{\theta}$ is the diagonal matrix formed by $\cos \theta_1, \dots, \cos \theta_\varrho$ along the diagonal. The $\cos \boldsymbol{\theta}$ is also known as the principal correlations or the canonical correlations ([Hotelling, 1936](#); [Golub and Zha, 1995](#)).

2.3 Metrics on Grassmannian Manifolds

In the following, we will review some of the most known Grassmannian metrics. Even if similar reviews could be found in [Edelman et al. \(1999\)](#) or in [Dhillon et al. \(2008\)](#), the one presented here is more detailed and more complete. Most of the Grassmannian metrics are relying on principal angles.

A metric, defined for an arbitrary non empty set X , is a function $d : X \times X \rightarrow \mathbb{R}^+ = [0, \infty)$ which verify these three axioms:

(A1) Separability $d(x, y) = 0$ iff $x = y$, $\forall x, y \in X$,

(A2) Symmetry $d(x, y) = d(y, x)$, $\forall x, y \in X$,

(A3) Triangle inequality $d(x, y) \leq d(x, z) + d(z, y)$, for all $x, y, z \in X$.

Let $\mathbf{Gr}(\varrho, n)$ be a Grassmannian manifold, as defined in Appendix A. Let \mathcal{U}, \mathcal{W} be two elements of $\mathbf{Gr}(\varrho, n)$ and let $\{\theta_1, \dots, \theta_\varrho\}$ be their associated principal angles.

The metrics on Grassmannian manifolds rely on principal angles and a general remark applying to all metrics based on principal angles is that they are significant if and only if \mathcal{U} and \mathcal{W} are defined in the same space \mathcal{V} . Furthermore, if \mathcal{U} and \mathcal{W} share basis vectors, an important proportion of the principal angles could be null. Assuming that U and W , such that $\text{span}(U) = \mathcal{U}$ and $\text{span}(W) = \mathcal{W}$, forms separated bases the maximum number of non-null angles is $\min(\varrho, n - \varrho)$.

1. Geodesic distance. A first metric to consider is the geodesic distance or arc length first introduced by [Wong \(1967\)](#).

Definition 2 Let \mathcal{U} and \mathcal{W} be two subspaces and $\boldsymbol{\theta} = [\theta_1, \dots, \theta_\varrho]$ be the ϱ principal angles between \mathcal{U} and \mathcal{W} , where ϱ is the dimension of the smallest subspace. The geodesic distance between \mathcal{U} and \mathcal{W} is:

$$d_g(\mathcal{U}, \mathcal{W}) = \|\boldsymbol{\theta}\|_2 = \left(\sum_{k=1}^{\varrho} \theta_k^2 \right)^{\frac{1}{2}}. \quad (1)$$

Nonetheless this distance is not differentiable everywhere ([Conway et al., 1996](#)). It takes values between zero and $\frac{\pi\sqrt{\varrho}}{2}$.

2. Chordal distance. Probably the most known Grassmannian metric ([Golub and van Loan, 1996](#); [Conway et al., 1996](#); [Wang et al., 2006](#)).

Definition 3 Let \mathcal{U} and \mathcal{W} be two subspaces and $\boldsymbol{\theta} = [\theta_1, \dots, \theta_\varrho]$ be the ϱ principal angles vector between \mathcal{U} and \mathcal{W} , where ϱ is the dimension of the smallest subspace. The chordal distance is defined as:

$$d_c(\mathcal{U}, \mathcal{W}) = \|\sin \boldsymbol{\theta}\|_2 = \left(\sum_{k=1}^{\varrho} \sin^2 \theta_k \right)^{\frac{1}{2}}. \quad (2)$$

Proposition 4 ([Dhillon et al., 2008](#)) The chordal distance can be reformulated in a more computationally effective way as:

$$d_c(\mathcal{U}, \mathcal{W}) = (\varrho - \|\underline{U}^T \underline{W}\|_F^2)^{\frac{1}{2}},$$

where the columns of \underline{U} and \underline{W} are the orthonormal bases for \mathcal{U} and \mathcal{W} , that is $\underline{U}^T \underline{U} = I_\varrho$ and $\text{span}(\underline{U}) = \mathcal{U}$.

The chordal distance approximates the geodesic distance when the planes are close and its square is differentiable everywhere.

The chordal distance could be seen as the embedding of an element \mathcal{U} of $\mathbf{Gr}(\varrho, n)$ with its projection matrix $\underline{U}\underline{U}^T$. The embedding could be formalized as:

$$\begin{aligned}\Pi_c : \mathbf{Gr}(\varrho, n) &\rightarrow \mathbb{R}^{n \times n} \\ \mathcal{U} &\mapsto \underline{U}\underline{U}^T .\end{aligned}$$

Thus the Grassmannian element is embedded in the set of orthogonal projection matrices of rank ϱ . The chordal distance could thus be expressed as:

$$d_c(\mathcal{U}, \mathcal{W}) = \frac{1}{\sqrt{2}} \|\underline{U}\underline{U}^T - \underline{W}\underline{W}^T\|_F .$$

Instead of using the Frobenius norm on the embedded Grassmannian, it is possible to use the 2-norm. The chordal 2-norm distance writes as:

$$d_{c2}(\mathcal{U}, \mathcal{W}) = \|\underline{U}\underline{U}^T - \underline{W}\underline{W}^T\|_2 = \|\sin \boldsymbol{\theta}\|_\infty = (\sin^2 \theta_\varrho)^{\frac{1}{2}} ,$$

which fails to be a metric because of (A1), it is thus a pseudo-metric.

A variation of the chordal metric is the projection metric, which starts with different assumptions but yield very similar results. The projection metric is the result of embedding $\mathbf{Gr}(\varrho, n)$ in the vector space $\mathbb{R}^{n \times \varrho}$:

$$\begin{aligned}d_{\text{proj}}(\mathcal{U}, \mathcal{W}) &= \|\underline{U}\underline{Y} - \underline{W}\underline{Z}\|_F \\ &= \left\| 2 \sin \left(\frac{\boldsymbol{\theta}}{2} \right) \right\|_2 = 2 \left(\sum_{k=1}^{\varrho} \sin^2 \frac{\theta_k}{2} \right)^{\frac{1}{2}} .\end{aligned}$$

The projection distance is very similar to the chordal one: it is the minimum distance between all the possible representations of the two subspaces \mathcal{U} and \mathcal{W} (Chikuse, 2003), which could be reformulated in term of the following minimization problem:

$$\begin{aligned}d_{\text{proj}}(\mathcal{U}, \mathcal{W}) &= \min_{\underline{R}_1, \underline{R}_2 \in \mathbf{O}(\varrho)} \|\underline{U}\underline{R}_1 - \underline{W}\underline{R}_2\|_F \\ &= \|\underline{U}\underline{Y} - \underline{W}\underline{Z}\|_F .\end{aligned}$$

where $\mathbf{O}(\varrho)$ is the orthogonal group.

As for the chordal distance, it is possible to rely on the 2-norm instead of the Frobenius norm for the projection metric, giving also the following pseudo-metric:

$$\begin{aligned}d_{\text{proj}2}(\mathcal{U}, \mathcal{W}) &= \|\underline{U}\underline{Y} - \underline{W}\underline{Z}\|_2 \\ &= \left\| 2 \sin \frac{\boldsymbol{\theta}}{2} \right\|_\infty = 2 \sin^2 \left(\frac{\theta_\varrho}{2} \right) .\end{aligned}$$

3. Fubini-Study distance. The Fubini-Study distance is derived via the Plücker embedding of $\mathbf{Gr}(\varrho, n)$ into the projective space $\mathbf{P}(\wedge^\varrho(\mathbb{R}^n))$ (taking the wedge product between all elements of \mathcal{W}), then taking the Fubini-Study metric (Kobayashi and Nomizu, 1969).

Definition 5 Let \mathcal{U} and \mathcal{W} be two subspaces and $\boldsymbol{\theta} = [\theta_1, \dots, \theta_\varrho]$ be the ϱ principal angles vector between \mathcal{U} and \mathcal{W} , where ϱ is the dimension of the smallest subspace. The Fubini-Study distance is defined as:

$$d_{\text{fs}}(\mathcal{U}, \mathcal{W}) = \arccos \left(\prod_{k=1}^{\varrho} \theta_k \right). \quad (3)$$

The Fubini-Study distance can be computed efficiently as:

$$d_{\text{fs}}(\mathcal{U}, \mathcal{W}) = \arccos(\det(\underline{U}^T \underline{W})) .$$

In Conway et al. (1996), the authors also argued that the Plücker embedding does not give a way to realize a geodesic or a chordal distance as an Euclidean distance.

4. Spectral distance. Introduced in Dhillon et al. (2008), the spectral distance is used to promote specific configurations of subspaces in packing problems.

Definition 6 Let \mathcal{U} and \mathcal{W} be two subspaces and $\boldsymbol{\theta} = [\theta_1, \dots, \theta_\varrho]$ be the ϱ principal angles vector between \mathcal{U} and \mathcal{W} , where ϱ is the dimension of the smallest subspace. The spectral distance is defined as:

$$\begin{aligned} d_{\text{s}}(\mathcal{U}, \mathcal{W}) &= \min_k \sin \theta_k \\ &= \left(1 - \|\underline{U}^T \underline{W}\|_{2,2}^2 \right)^{\frac{1}{2}}, \end{aligned} \quad (4)$$

where the spectral norm $\|X\|_{2,2}$ returns the largest singular value of X .

This distance allows to solve the packing problem with equi-isoclinic configurations, that is when all principal angles formed by all pairs of subspaces are equal. Nonetheless this distance, defined in Equation (4), is not a metric.

5. Binet-Cauchy distance. A last metric on $\mathbf{Gr}(\varrho, n)$ is the Binet-Cauchy distance which is the product of principal angles or canonical correlations (Wolf and Shashua, 2003; Vishwanathan and Smola, 2004).

Definition 7 Let \mathcal{U} and \mathcal{W} be two subspaces and $\boldsymbol{\theta} = [\theta_1, \dots, \theta_\varrho]$ be the ϱ principal angles vector between \mathcal{U} and \mathcal{W} , where ϱ is the dimension of the smallest subspace. The Binet-Cauchy distance is defined as:

$$d_{\text{bc}}(\mathcal{U}, \mathcal{W}) = \left(1 - \prod_{k=1}^{\varrho} \cos^2 \theta_k \right)^{\frac{1}{2}}. \quad (5)$$

Distance	Definition	Metric	DE
Geodesic	$d_g(\mathcal{U}, \mathcal{W}) = \ \boldsymbol{\theta}\ _2$	✓	×
Chordal	$d_c(\mathcal{U}, \mathcal{W}) = \ \sin \boldsymbol{\theta}\ _2$	✓	✓
Fubini-Study	$d_{fs}(\mathcal{U}, \mathcal{W}) = \arccos(\det(\underline{U}^T \underline{W}))$	✓	✓
Spectral	$d_s(\mathcal{U}, \mathcal{W}) = (1 - \ \underline{U}^T \underline{W}\ _{2,2}^2)^{\frac{1}{2}}$	×	×
Binet-Cauchy	$d_{bc}(\mathcal{U}, \mathcal{W}) = (1 - \prod_k \cos^2 \theta_k)^{\frac{1}{2}}$	✓	✓

Table 1: Distances on Grassmannian manifolds \mathcal{U} and \mathcal{W} , DE stands for differentiable everywhere.

This distance was introduced to enhance the computational efficiency and the numerical stability of the Canonical Correlation Analysis (CCA) based kernel approaches.

All these distances are summarized in the Table 1 with their definition.

3 Frames

We assume that the reader is familiar with frame theory and we restrict ourselves to introduce definitions and facts that are mandatory for further developments in the paper. Please refer to [Christensen \(2003\)](#) for a deeper study of the topic.

We assume that \mathcal{V} , together with its inner product and the induced norm, is a separable Hilbert space. The indexing set I is in \mathbb{N} as we will consider, without loss of generality, only finite-dimensional frames.

Definition 8 *A family of elements $U = \{u_i\}_{i \in I}$ is a frame for \mathcal{V} if there exist $0 < b_1 \leq b_2 \leq \infty$, such that for all $w \in \mathcal{V}$:*

$$b_1 \|w\|^2 \leq \sum_{i \in I} |\langle w, u_i \rangle|^2 \leq b_2 \|w\|^2 . \quad (6)$$

If $\|u_i\|^2 = 1$ for all i , it is a unit norm frame. A tight frame is a frame whose bounds are equal ($b_1 = b_2$) and if a frame verifies $b_1 = b_2 = 1$, it is a Parseval frame or a normalized tight frame. This inequality implies that the frame U is full rank ([Mallat, 1999](#)).

Definition 9 *Let $\{u_i\}_{i \in I}$ be a frame for \mathcal{V} . The Bessel map, or analysis operator, is a function $T : \mathcal{V} \rightarrow \ell^2(I)$ defined by:*

$$T(w) = (\langle w, u_i \rangle)_{i \in I} .$$

Definition 10 Let $\{u_i\}_{i \in I}$ be a frame for \mathcal{V} . The pre-frame operator, or the synthesis operator, is a map $T^* : L^2(I) \rightarrow \mathcal{V}$ defined by:

$$T^*c = \sum_{i \in I} c_i u_i .$$

For an orthonormal basis, these operators are invertible and are each other inverse. In the context of frames, it is usually not the case. Specifically, T^* is always surjective but usually not injective and T is always injective but usually not surjective.

Definition 11 The frame operator is a mapping $S : \mathcal{V} \rightarrow \mathcal{V}$ and is defined as T^*T , that is

$$S(w) = \sum_{i \in I} \langle w, u_i \rangle u_i .$$

The frame operator is well defined, bounded, self-adjoint and positive.

4 Sparse Coding, Dictionary Learning and Grassmannian Frames

Our attempt here is to establish a principled connexion between three close but still disjoint fields of research, namely Grassmannian packing, dictionary learning and compressed sensing. In this section, without loss of generality, we will restrict the discussion to the case of $\rho = 1$.

4.1 Packing Problem and Grassmannian Frames

One way to characterize a frame is to evaluate how “overcomplete” it is. The redundancy expresses this measure of overcompleteness and is defined for $U = \{u_i\}_{i=1}^m$ in a n -dimensional vector space \mathcal{V} as $\frac{m}{n}$. A more informative measure is the coherence (or maximum correlation), defined as the maximum of the absolute inner product between any two elements of a unit norm frame $\{u_i\}_{i=1}^m$:

$$\mu = \max_{i,j \in \{1, \dots, m\}, i \neq j} \{ |\langle u_i, u_j \rangle| \} . \quad (7)$$

The constraint on unit norm frame could be removed if the inner product of Equation (7) is divided by norm of the frame elements.

A fundamental result in frame theory bounds for the coherence of any frame and the bound depends only on the number of elements in the frame and the dimension of the associated space \mathcal{V} .

Theorem 12 (Welch, 1974) Let $U = \{u_i\}_{i=1}^m$ be a frame for \mathcal{V} , then its Welch bound is

$$\mu \geq \sqrt{\frac{m-n}{n(m-1)}} . \quad (8)$$

The equality hold iff U is an Equiangular Tight Frame (ETF), $|\langle u_i, u_j \rangle| = c$ for all i, j with $i \neq j$, for some constant $c \geq 0$. Furthermore, the equality holds only if $m \leq \frac{n(n+1)}{2}$.

In order to construct equiangular tight frames, a possible approach is to define a set of frames which minimizes the coherence for all unit norm frames with the same redundancy.

Definition 13 (Strohmer and Heath Jr., 2003) A sequence of vectors $U = \{u_i\}_{i=1}^m$ in \mathcal{V} is called a Grassmannian frame if it is the solution to $\min \mu$ where the minimum is taken over all unit norm frames in \mathcal{V} .

A unit norm frame holding the equality in Equation (8) is called an optimal Grassmannian frame. Note that optimal Grassmannian frames do not necessarily exist for all packing problem: a known example (Conway et al., 1996) is the packing of 5 vectors in \mathbb{R}^3 , with $\varrho = 1$, $m = 5$, $n = 3$, whose coherence is at best $\frac{1}{\sqrt{5}}$ while its Welch bound is $\frac{1}{\sqrt{6}}$.

The ETF and the Grassmannian frames offer a well suited tool to solve the packing problem. In its general formulation, for given n , m and ϱ , the goal of the packing problem is to find a set of ϱ -dimensional subspaces \mathcal{U}_i in a n -dimensional vector space \mathcal{V} such that (i) the minimal distance between any two of these subspaces are as large as possible, or equivalently (ii) the maximal correlation between subspaces is as small as possible. A schematic view of the packing of a space is shown on left part of Figure 2.

In the case of $\mathbf{Gr}(1, n)$, investigated in Strohmer and Heath Jr. (2003), the correlation depends only on n and m . Thus subspaces are lines passing through the origin in \mathcal{V} and the angles between any two of the m lines should be as large as possible. As maximizing the angle between lines is the same as minimizing the absolute value of the inner product of the u_i vectors, finding an optimal packing in $\mathbf{Gr}(1, n)$ is equivalent to finding a Grassmannian frame.

4.2 Compressed Sensing vs Grassmannian Frames

It is interesting to see that the ETF are also linked with sparse coding and compressed sensing. Very briefly, sparse representations and compressed sensing (Donoho, 2006) proposes to acquire a signal y in \mathcal{V} by collecting m linear measurements of the form $a_k = \langle u_k, y \rangle + e_k$ with $1 \leq k \leq m$ or equivalently:

$$a = U^T y + e ,$$

where U is a $n \times m$ sensing matrix, with $m \ll n$ and e is an error term. The compressed sensing theory states that, if y is reasonably sparse, it is possible to recover y by convex programming, as it is the solution to

$$\min_{\tilde{y} \in \mathcal{V}} \|\tilde{y}\|_1 \quad \text{s.t.} \quad \|a - U^T \tilde{y}\|_2 \leq \epsilon . \quad (9)$$

In a noiseless settings, with a null vector e , to achieve an exact reconstruction of y , a sensing matrix should satisfy the (K, δ) -Restricted Isometry Property (RIP).

Definition 14 For a $n \times m$ sensing matrix U , the K -restricted isometry constant δ of U is the smallest quantity such that

$$(1 - \delta) \|y\|_2^2 \leq \|U^T y\|_2^2 \leq (1 + \delta) \|y\|_2^2$$

holds for every K -sparse vectors y .

A matrix with a small δ constant indicates that every subset of K or less columns is linearly independent. Furthermore, any U sensing matrix which is $(2K, \delta)$ -RIP for some $\delta < \sqrt{2} - 1$ allows to recover the signal using Equation (9) with the following error bound

$$\|\tilde{y} - y\|_2 \leq C_0 \frac{\|\tilde{y} - y\|_1}{\sqrt{K}} + C_1 \epsilon .$$

Matrices with small δ are abundant and could be easily obtained. Matrices with Gaussian or Bernoulli entries have small δ with very high probability if the number of measurement m is of the order of $K \log(n/K)$. However, testing a matrix for RIP is computationally intensive, as it requires to determine singular values for n -by- K submatrices.

An alternative approach is to rely on ETF for designing deterministically RIP matrices. From the Theorem 2 of [Mixon et al. \(2011\)](#), the following lemma are defined:

Lemma 15 ([Mixon et al., 2011](#)) Let $U = \{u_i\}_{i \in I}$ be a unit norm frame on \mathcal{V} . The smallest δ_{\min} for which U is (K, δ) -RIP is defined by

$$\delta_{\min} = \max_{\mathcal{S} \subseteq \{1, \dots, n\}, |\mathcal{S}|=K} \|U_{\mathcal{S}}^T U_{\mathcal{S}} - I_{\mathcal{S}}\|_2 ,$$

where $U_{\mathcal{S}}$ is the submatrix consisting of columns of U indexed by \mathcal{S} .

As it is computationally demanding to compute all the eigenvalues of $\mathcal{S} \times \mathcal{S}$ submatrices, the Gershgorin circle theorem is used to obtain a coarse bound linked to RIP ([DeVore, 2007](#); [Applebaum et al., 2009](#)).

Lemma 16 ([Mixon et al., 2011](#)) Let $U = \{u_i\}_{i \in I}$ be a unit norm frame, then the K -restricted isometry constant

$$\delta \leq (K - 1)\mu ,$$

where μ is the coherence.

Thus, U is (K, δ) -RIP for $\delta \geq (K - 1)\mu$ and using Welch bound given in Equation (8). As ETF meet the Welch bound with equality, they are (K, δ) for $\delta^2 \geq \frac{(K-1)^2(n-m)}{m(n-1)}$. Hence, ETF are good candidates for producing deterministically RIP sensing matrices.

4.3 Dictionary Learning vs Grassmannian Frames

The dictionary learning aims at capturing most of the energy of a signal database $Y = \{y_j\}_{j=1}^q \subset \mathcal{V}$ and representing it through a collection $U = \{u_i\}_{i=1}^m$ thanks to a set of sparse coefficients $A = \{a_j\}_{j=1}^q$. This collection, which is redundant since $m \gg n$, is called overcomplete dictionary. This dictionary is not necessarily full rank, which is a difference from frames as imposed by Equation (6). Formally, the dictionary learning problem writes as:

$$\min_{U,A} \|Y - UA\|_F^2 \quad \text{s.t.} \quad \|a_j\|_0 \leq K, j = 1, \dots, q \\ \text{and} \quad \|u_i\|_2 = 1, i = 1, \dots, m,$$

where $\|a_j\|_0$ is the number of nonzero elements of vector a_j . This problem is tackled by dictionary learning algorithms (DLAs), in which energy representative patterns of the dataset are iteratively selected by a sparse approximation step, and then updated by a dictionary update step. Different algorithms deal with this problem: see for example the method of optimal directions proposed by Engan et al. (2000) and generalized under the name iterative least-squares DLA (Engan et al., 2007), the K-SVD (Aharon et al., 2006) or the online DLA (Mairal et al., 2010).

Data-driven dictionary learning extracts atoms that endow most of the energy of the dataset in the underlying norm sense. They are learned empirically to be the optimal dictionary that jointly gives sparse approximations for all of the signals of this set (Tošić and Frossard, 2011a), as illustrated on the center part of Figure 2. During learning process, no constraint is put on the dictionary elements in term of coherence. Whereas in the packing procedure for m elements, as shown in Section 4.1, amounts to come up with elements with maximal distance between any two of them, hence minimizing the maximal correlation.

The packing procedure is data-free, in the sense that it is independent of the data and is a pure geometrical problem, while the dictionary learning is data-driven. An interesting view is to connect the two visions, an approach which is emerging in recent problem. A possible approach is to constraint the packing problem on the energy of the dataset, in other words a data-constraint packing.

An alternative point of view, as developed by Yaghoobi et al. (2009), is to constraint the dictionary learning with a penalty term on the coherence of the dictionary elements (see the right part of Figure 2). This could be formalized as adding the penalty term during the update of dictionary, to constraint the dictionary element to be close to ETF, as:

$$\|U^T U - G\|_F^2 \leq \epsilon \quad \text{with} \quad G = \begin{bmatrix} 1 & & g_{ij} \\ & \ddots & \\ g_{ji} & & 1 \end{bmatrix},$$

where $\max_{i \neq j} |g_{ij}|$ is less or equal to the Welch bound given in Equation (8).

As a summary, and as illustrated on Figure 2, packing spans the data space equi-correlatively (left), dictionary learning spans the data energy sparsely (center), and coherence constrained dictionary learning spans the data energy sparsely and equi-correlatively (right).




Figure 2: In the frame identification problem or packing problem (left), the frame elements (circles) are used cover the entire space under the constraint that the distance between any two elements should be maximal, being as close as possible to an ETF. In the dictionary learning setup (center), the frame elements are chosen to minimize the reconstruction error on the energy of the dataset (represented as gray values in “G” shape). An hybrid approach (right) adds coherence constraints in the dictionary learning framework.

5 Metrics between multivariate dictionaries

In this section, we will use the metrics described in Section 2.3 and summarized in Table 1 which act on elements of a Grassmannian as ground distance. This ground distance has a key role for the definition of a metric between sets of points in a Grassmannian space, that is a distance between two subspaces spanned by their associated dictionaries.

5.1 Preliminaries

In this section, the vector space \mathcal{V} is of dimension $n \times \rho$, its elements will be denoted as $U, W \in \mathcal{V}$. Their respective spans are $\text{span}(U) = \mathcal{U}$ and $\text{span}(W) = \mathcal{W}$. The indexed families of elements will be denoted $\mathbf{U} = \{U_i\}_{i \in I}$ and $\mathbf{W} = \{W_j\}_{j \in J}$. Indexed families of subspaces will be denoted $\mathbb{U} = \{\mathcal{U}_i\}_{i \in I}$ and $\mathbb{W} = \{\mathcal{W}_j\}_{j \in J}$.

The Grassmannian manifold $\mathbf{Gr}(\varrho, n)$, $\varrho \leq \rho$, together with any of the distances d discussed in Section 2.3, defines a metric space (or pseudo-metric space depending on the properties of the underlying distance). We will

denote it in the sequel as (\mathbb{G}, d) , and when there is no confusion as \mathbb{G} . The following theorem characterizes the properties of Grassmannian manifolds.

Theorem 17 (Milnor and Stasheff, 1974) $\text{Gr}(\varrho, n)$ is a Hausdorff, compact, connected smooth manifold of dimension $\varrho(n - \varrho)$.

Thus, we know that \mathbb{G} is complete and (separable) Hausdorff, and from the fact that any topological manifold is locally compact, then one can define a Borel measure on \mathbb{G} . We will denote by π such a measure on the Borel σ -algebra of \mathbb{G} . The support of π is denoted $\text{Spt}(\pi)$ and is the smallest closed set on which it is concentrated, that is the minimal closed subset $X \subset \mathbb{G}$ such that $\pi(\mathbb{G} \setminus X) = 0$. The triplet (\mathbb{G}, d, π) is then a metric measure space. Furthermore, \mathbb{G} is a Polish space, that is a separable completely metrizable topological space.

5.2 Hausdorff Distance

A classical approach to define a distance between subsets of a metric space is to rely on the Hausdorff distance. Within our context, frames are subsets of the space \mathbb{G} , and the Hausdorff distance is a first natural way to define a distance in the frame space.

We denote by $B_r(\mathbb{W})$ the r -neighborhood of a set \mathbb{W} in the metric space (\mathbb{G}, d) , that is the set of points \mathcal{U} such that $\inf\{d(\mathcal{U}, \mathcal{W}) : \mathcal{W} \in \mathbb{W}\} < r$.

Definition 18 Let $\mathbb{U} = \{\mathcal{U}_i\}_{i \in I}$ and $\mathbb{W} = \{\mathcal{W}_j\}_{j \in J}$ be two subsets of the metric space (\mathbb{G}, d) . The Hausdorff distance between \mathbb{U} and \mathbb{W} , denoted by $d_H(\mathbb{U}, \mathbb{W})$, is defined as

$$d_H(\mathbb{U}, \mathbb{W}) := \inf\{r > 0 : \mathbb{U} \subset B_r(\mathbb{W}) \text{ and } \mathbb{W} \subset B_r(\mathbb{U})\} ,$$

or equivalently as:

$$d_H(\mathbb{U}, \mathbb{W}) := \max \left(\sup_{\mathcal{U} \in \mathbb{U}} \inf_{\mathcal{W} \in \mathbb{W}} d(\mathcal{U}, \mathcal{W}), \sup_{\mathcal{W} \in \mathbb{W}} \inf_{\mathcal{U} \in \mathbb{U}} d(\mathcal{U}, \mathcal{W}) \right) .$$

Proposition 19 (Burago et al., 2001) Consider the metric space \mathbb{G} , then:

1. d_H is a semi-metric on $2^{\mathbb{G}}$ (the set of all subsets of \mathbb{G}) ,
2. $d_H(\mathbb{U}, \hat{\mathbb{U}}) = 0$ for any $\mathbb{U} \subset \mathbb{G}$ where $\hat{\mathbb{U}}$ denotes the closure of \mathbb{U} ,
3. If \mathbb{U} and \mathbb{W} are closed subsets of \mathbb{G} and $d_H(\mathbb{U}, \mathbb{W}) = 0$, then $\mathbb{U} = \mathbb{W}$.

Then it turns that the Hausdorff distance is a true metric iff the considered sets are closed. A known limitation is that the Hausdorff distance is non-smooth.

This distance could be reformulated in term of a *correspondence*. This allows to explicit the link between the Hausdorff distance and the class of Wasserstein distances that will be introduced in the sequel (cf. Section 5.3).

Definition 20 Let \mathbb{U} and \mathbb{W} be two sets, a subset $R \subset \mathbb{U} \times \mathbb{W}$ is a correspondence if $\forall \mathcal{U} \in \mathbb{U}$ there is $\mathcal{W} \in \mathbb{W}$ s.t. $(\mathcal{U}, \mathcal{W}) \in R$ and $\forall \mathcal{W} \in \mathbb{W}$ there is $\mathcal{U} \in \mathbb{U}$ s.t. $(\mathcal{U}, \mathcal{W}) \in R$.

Let $\mathcal{R}(\mathbb{U}, \mathbb{W})$ denotes the set of all possible correspondences between \mathbb{U} and \mathbb{W} . It is possible to distinguish between a correspondence ϕ between \mathbb{U} and \mathbb{W} defined on the subset $R \subset \mathbb{U} \times \mathbb{W}$ as

$$R = \{(\mathcal{U}, \mathcal{W}) \in \mathbb{U} \times \mathbb{W} : \mathcal{W} \in \phi(\mathcal{U}), \mathcal{U} \in \phi(\mathcal{W})\}$$

and the subset R , which is also called the graph of ϕ .

Proposition 21 Let $\mathbb{U} = \{\mathcal{U}_i\}_{i \in I}$ and $\mathbb{W} = \{\mathcal{W}_j\}_{j \in J}$ be two subsets of the metric space (\mathbb{G}, d) . The Hausdorff distance between \mathbb{U} and \mathbb{W} could be rewritten as:

$$d_H(\mathbb{U}, \mathbb{W}) = \inf \left\{ \sup_{\mathcal{U}, \mathcal{W} \in R} d(\mathcal{U}, \mathcal{W}) : R \in \mathcal{R}(\mathbb{U}, \mathbb{W}) \right\},$$

or equivalently

$$d_H(\mathbb{U}, \mathbb{W}) = \inf_{R \in \mathcal{R}(\mathbb{U}, \mathbb{W})} \sup_{\mathcal{U}, \mathcal{W} \in R} d(\mathcal{U}, \mathcal{W}). \quad (10)$$

With the view of this last proposition, one can easily construct new set-metrics by replacing the sup operator in Equation (10) by any ℓ_p norm, hence leading to more interesting smooth set-metrics. Such metrics belong to the class of Wasserstein distances by restricting the Borel measure to the Dirac function. The next section clarifies this statement.

5.3 Wasserstein Distance

Let us define by $\mathcal{C}(\mathbb{G})$ any collection on \mathbb{G} and by $\mathbb{G}_W = \{(\mathcal{U}, \pi_{\mathcal{U}}) : \mathcal{U} \in \mathbb{G}\}$ where $\pi_{\mathcal{U}}$ is a Borel measure.

Definition 22 Let $\mathbb{U}, \mathbb{W} \in \mathbb{G}_W$. A measure π on the product space $\mathbb{U} \times \mathbb{W}$ is a coupling of $\pi_{\mathbb{U}}$ and $\pi_{\mathbb{W}}$ if

$$\pi(\bar{\mathbb{U}} \times \mathbb{W}) = \pi_{\mathbb{U}}(\bar{\mathbb{U}}), \quad \pi(\mathbb{U} \times \bar{\mathbb{W}}) = \pi_{\mathbb{W}}(\bar{\mathbb{W}}) \quad (11)$$

for all Borel sets $\bar{\mathbb{U}} \subset \mathbb{U}$, $\bar{\mathbb{W}} \subset \mathbb{W}$.

We denote by $\mathcal{M}(\pi_{\mathbb{U}}, \pi_{\mathbb{W}})$ the set of all couplings of $\pi_{\mathbb{U}}$ and $\pi_{\mathbb{W}}$. An interesting property is that given $\pi \in \mathcal{M}(\pi_{\mathbb{U}}, \pi_{\mathbb{W}})$ then $R(\pi) = \text{Spt}(\pi)$ is in $\mathcal{R}(\mathbb{U}, \mathbb{W})$.

Definition 23 Let $\mathbb{U}, \mathbb{W} \in \mathbb{G}_W$, and π a coupling as defined in Equation (11). Given $p > 0$, the Wasserstein distance is defined as:

$$d_W^p(\mathbb{U}, \mathbb{W}) = \inf_{\pi \in \mathcal{M}(\pi_{\mathbb{U}}, \pi_{\mathbb{W}})} \left(\int_{\mathbb{U} \times \mathbb{W}} d(\mathcal{U}, \mathcal{W})^p d\pi(\mathcal{U}, \mathcal{W}) \right)^{\frac{1}{p}} \quad \text{if } 1 \leq p \quad (12)$$

$$d_W^p(\mathbb{U}, \mathbb{W}) = \inf_{\pi \in \mathcal{M}(\pi_{\mathbb{U}}, \pi_{\mathbb{W}})} \int_{\mathbb{U} \times \mathbb{W}} d(\mathcal{U}, \mathcal{W})^p d\pi(\mathcal{U}, \mathcal{W}) \quad \text{if } 0 < p < 1 \quad (13)$$

$$d_W^\infty(\mathbb{U}, \mathbb{W}) = \inf_{\pi \in \mathcal{M}(\pi_{\mathbb{U}}, \pi_{\mathbb{W}})} \sup_{(\mathcal{U}, \mathcal{W}) \in \mathcal{R}(\pi)} d(\mathcal{U}, \mathcal{W}) \quad \text{otherwise} . \quad (14)$$

This distance is also called the Wasserstein-Kantorovich-Rubinstein.

Proposition 24 Let \mathbb{U}, \mathbb{W} be two subsets of \mathbb{G} , if we consider as a measure π the Dirac one, then the Wasserstein distance is related to the Hausdorff distance by the following inequality:

$$d_H(\mathbb{U}, \mathbb{W}) \leq d_W^\infty((\mathbb{U}, \pi_{\mathbb{U}}), (\mathbb{W}, \pi_{\mathbb{W}})) . \quad (15)$$

5.4 Set-Metrics for Frames and dictionaries

Metrics in Equations (12), (13), (14) and (15) are defined on Grassmannian space. As indicated in Section 5.1, (\mathbb{G}, d) is a separable metric space allowing to compute a distance between the collection of subspaces spanned by a frame and the collection of subspaces spanned by another. The distance between two frames in the frame space, that is \mathcal{V} , is defined as follows.

Definition 25 Let $\mathbf{U} = \{U_i\}_{i \in I}$ and $\mathbf{W} = \{W_j\}_{j \in J}$ be two frames of the $n \times \rho$ vector space \mathcal{V} . We define a distance between these two frames as:

$$d_F(\mathbf{U}, \mathbf{W}) = d_S(\mathbb{U}, \mathbb{W}) , \quad (16)$$

where $\mathbb{U} = \{\mathcal{U}_i : \mathcal{U}_i = \text{span}(U_i), i \in I\}$ and $\mathbb{W} = \{\mathcal{W}_j : \mathcal{W}_j = \text{span}(W_j), j \in J\}$, and d_S is either d_H or d_W .

From this definition, we can formulate the proposition:

Proposition 26 Let $\mathcal{C}(\mathcal{V})$ be any collection on \mathcal{V} . Then the following holds:

- d_F is pseudo-metric and hence $(\mathcal{C}(\mathcal{V}), d_F)$ is pseudo-metric space,
- d_F is invariant by linear combinations.

Proof The proof is a direct consequence of Equation (16) since d_F is defined as a distance between subspaces. ■

The frame distance d_F is defined as the distance between their subspaces, that is d_F is acting in the frame space \mathcal{V} whereas the distance d_S is acting

in the Grassmannian space \mathbb{G} . As an element \mathcal{U} of \mathbb{G} is a subspace, there exist an infinite number of elements U in \mathcal{V} spanning \mathcal{U} . Thus two distinct frames $\mathbf{U}_1 \neq \mathbf{U}_2$, that is two collections $\{U_i^1\}_{i \in I}$ and $\{U_i^2\}_{i \in I}$ of elements in \mathcal{V} , could span the same collection of subspaces $\mathbb{U} = \{\mathcal{U}_i\}_{i \in I}$ in \mathbb{G} . In other words, a distance $d_F(\mathbf{U}_1, \mathbf{U}_2) = d_S(\mathbb{U}, \mathbb{U}) = 0$ could exist for two separate frames $\mathbf{U}_1 \neq \mathbf{U}_2$. As the separability axiom does not hold, d_F is a pseudo-metric and the separability axiom is relaxed to the identity axiom: $d(x, x) = 0, \forall x \in X$.

Pseudo-metric spaces are similar to metric spaces and important properties of metric spaces could be proven as well for pseudo-metric spaces, so the fact that $(\mathcal{C}(\mathcal{V}), d_F)$ is a pseudo-metric space has a very limited impact for further theoretical developments. This situation is even desirable, as the distance between frames is defined through their associated subspaces yields the invariance by linear combinations, which is an essential property for comparing multivariate dictionaries.

The proposition of set-metrics for frames is a major contribution for the dictionary learning community as it is the first proposition to define a real metric in the dictionary space. Previous works rely on heuristics to assess differences between dictionaries and no tractable framework has been proposed to compare dictionaries. The work of [Skretting and Engan \(2011\)](#) is a noticeable attempt to quantify the similarity between two dictionaries, without achieving to define a metric.

The proposed pseudo-metrics – relying on Hausdorff and Wasserstein set-metrics, as well as chordal, Fubini-Study or other ground metrics – have been described in the literature and are easy to implement. The following section described a multivariate dictionary learning algorithm which is then applied in Section 7 on synthetic data and in Section 8 on real world signals, demonstrating that these metrics could be applied “out-of-the-box” to design better dictionary learning algorithms or to assess machine learning datasets.

6 Dictionary Learning for Multivariate Signals

Hereafter, to illustrate the metrics between frames with concrete examples, we will consider a multivariate dictionary \mathbf{U} as a collection of m multivariate atoms $U_i \in \mathbb{R}^{n \times \rho}$. In this section, models and dictionary learning algorithms for multicomponent signals are detailed.

6.1 Model for Multivariate Dictionaries

We consider a multivariate signal $Y \in \mathbb{R}^{n \times \rho}$, whose decomposition is carried out on the multivariate dictionary \mathbf{U} thanks to the multivariate model ([Barthélemy](#)

et al., 2012):

$$Y = \sum_{i=1}^m a_i U_i + E, \quad (17)$$

where $a_i \in \mathbb{R}$ denotes the coding coefficient and $E \in \mathbb{R}^{n \times \rho}$ the residual error. Two problems need to be handled on this model: the first one, called sparse approximation, estimates the coefficient vector a when the dictionary \mathbf{U} is fixed, and the second one, called dictionary learning, estimates the best dictionary \mathbf{U} from the dataset $\mathbf{Y} = \{Y_j\}_{j=1}^q$.

Since the dictionary is redundant, $m \gg n$, the linear system of (17) is thus under-determined and has multiple solutions. The introduction of constraints such as sparsity allows to regularize the solution. With the sparse approximation, only K active atoms are selected among the m possible ones and the associated coefficients vector a is computed to maximize the approximation of the signal Y . The multivariate sparse approximation can be formalized as:

$$\min_a \left\| Y - \sum_{i=1}^m a_i U_i \right\|_F^2 \quad \text{s.t.} \quad \|a\|_0 \leq K, \quad (18)$$

where $K \ll m$ is a constant. But this problem is NP-hard (Davis, 1994). In the univariate case, non-convex pursuits tackle it sequentially, such as matching pursuit (MP) (Mallat and Zhang, 1993) or orthogonal matching pursuit (OMP) (Pati et al., 1993). Many other sparse approximation algorithms are reviewed in Tropp and Wright (2010). In the multivariate case, Equation (18) is solved by the multivariate OMP (M-OMP) (Barthélemy et al., 2012).

Dictionary learning presented in Section 4.3 is now extended to the multivariate case. Adding an index j to variables to denote their link to the signal Y_j , the multivariate dictionary learning problem is formalized by:

$$\min_{\mathbf{U}} \sum_{j=1}^q \min_{a_j} \left\| Y_j - \sum_{i=1}^m a_{i,j} U_i \right\|_F^2 \quad \text{s.t.} \quad \|a_j\|_0 \leq K, \quad j = 1 \dots q \quad (19)$$

and $\|U_i\|_F = 1, \quad i = 1 \dots m.$

The multivariate DLA (M-DLA) proposed in Barthélemy et al. (2013a) solves this problem by learning a multivariate dictionary from the multivariate dataset \mathbf{Y} . This algorithm has been used to study multivariate time-series such as 2D spatial trajectories (Barthélemy et al., 2012) and EEG signals (Barthélemy et al., 2013a).

6.2 Model for Multicomponent Dictionaries

Some other works use multicomponent dictionaries to analyze multicomponent data: audio-visual data (Monaci et al., 2007, 2009), electro-cardiogram

signals (Mailhé et al., 2009), color images (Mairal et al., 2008), hyperspectral images (Moudden et al., 2009), stereo images (Tošić and Frossard, 2011b). But their decomposition model is different from the one used with multivariate dictionaries.

With multicomponent dictionaries, each component is associated with a univariate atom multiplied by its own coefficient, contrary to the multivariate model of Equation (17) where components form a multivariate atom multiplied by one unique coefficient, identical for all components. Let recall that $Y = [y_1, \dots, y_\rho]$ is a $n \times \rho$ matrix and $U_i = [u_{i,1}, \dots, u_{i,\rho}]$ is a $n \times \rho$ matrix. In this paragraph, the coefficients $a_i = [a_{i,1}, \dots, a_{i,\rho}]$ form a ρ -dimensional vector. Each component k is decomposed independently:

$$y_k = \sum_{i=1}^m a_{i,k} u_{i,k} + e_k, \quad k = 1 \dots \rho, \quad (20)$$

thus the coefficients a is a $m \times \rho$ matrix. To compare the multivariate model of Equation (17) with the multicomponent model of Equation (20), the latter can be reformulated as:

$$Y = \sum_{i=1}^m a_i \otimes U_i + E,$$

where \otimes is defined as the element-wise product along dimension ρ . Thus, in the multicomponent approach, multivariate signals are transformed in parallel univariate signals, causing an important information loss.

Sparse approximations adapted to this model estimate the multicomponent coefficients associated to the selected active atoms. The relation between the components of an atom is only considered during the atom selection: the index i is the same for the ρ components $u_{i,k}$. The selected atom U_i is thus jointly chosen between components thanks to a sparse mixed norm (Gribonval et al., 2007; Rakotomamonjy, 2011). Algorithms adapted to multicomponent approach are not detailed more here since this model is not used in the sequel.

6.3 Model for Rotation Invariance

The multivariate model described in Equation (17) is now extended to include rotation invariance. A rotation matrix $R_i \in \mathbb{R}^{\rho \times \rho}$ is added to each atom $U_i \in \mathbb{R}^{n \times \rho}$, with $R_i R_i^T = I_\rho$. Thus, the decomposition model becomes:

$$Y = \sum_{i=1}^m a_i U_i R_i + E. \quad (21)$$

This model is called *nD Rotation Invariant* (nDRI) in Barthélemy et al. (2013b), since multivariate atoms U_i are now able to rotate.

The multivariate sparse approximation described in Equation (18) is extended to this rotation invariant model. Coefficients $a = \{a_i\}_{i=1}^m$ and rotation matrices $\mathbf{R} = \{R_i\}_{i=1}^m$ are estimated with the nDRI-OMP (Barthélemy et al., 2013b) such as:

$$\min_{a, \mathbf{R}} \left\| Y - \sum_{i=1}^m a_i U_i R_i \right\|_F^2 \quad \text{s.t.} \quad \|a\|_0 \leq K$$

$$\text{and } R_i R_i^T = I_\rho, i = 1 \dots m .$$

The core of the nDRI-OMP is the orthogonal Procrustes problem. The goal is to estimate the coefficient $\alpha \in \mathbb{R}$ and the rotation matrix $R \in \mathbb{R}^{\rho \times \rho}$ which optimally register the atom $U \in \mathbb{R}^{n \times \rho}$ on the signal $Z \in \mathbb{R}^{n \times \rho}$:

$$(\tilde{\alpha}, \tilde{R}) = \arg \min_{\alpha, R} \| Z - \alpha U R \|_F^2 \quad \text{s.t.} \quad R R^T = I_\rho .$$

The function which computes this step is denoted $(\tilde{\alpha}, \tilde{R}) = \text{nD_Registration}(U, Z)$ in the following.

The multivariate dictionary learning described in Equation (19) is extended to the rotation invariant case as:

$$\min_{\mathbf{U}} \sum_{j=1}^q \min_{a_j, \mathbf{R}_j} \left\| Y_j - \sum_{i=1}^m a_{i,j} U_i R_{i,j} \right\|_F^2$$

$$\text{s.t.} \quad \|a_j\|_0 \leq K, j = 1 \dots q$$

$$\text{and } R_{i,j} R_{i,j}^T = I_\rho, i = 1 \dots m, j = 1 \dots q$$

$$\text{and } \|U_i\|_F = 1, i = 1 \dots m .$$

The algorithm which solves this problem is called nDRI-DLA (Barthélemy et al., 2013b) and it has been applied in a trivariate case to study 3D spatial trajectories.

7 Experiments on Synthetic Data

This work is the first attempt to assess qualitatively or quantitatively the dictionary learning algorithms with real metrics. Metrics on frames have several immediate and useful applications to improve our understanding of learning algorithms and datasets in a dictionary-based framework. This section is devoted to show why relying on metrics allows to dramatically improve the assessment of dictionary learning algorithms. More precisely, a set of experiments is conducted in Section 7.3 to reproduce state-of-the-art results on synthetic dataset and shows how the different proposed metrics compared with the commonly used indicators. A second set of experiments, presented in Section 8, shows how metrics could be used on multivariate signal, in the context of Brain-Computer Interface, to evaluate qualitatively datasets of brain signal.

7.1 Experimental Protocol

The experimental protocol described hereafter was inspired by Aharon (2006), Barthélemy et al. (2012) and Barthélemy et al. (2013b)¹. It is presented leaving out the explanations relative to the shift-invariant signal processing to lighten the description.

A dictionary with m multivariate atoms of $\rho = 10$ components, generated randomly, serve as a generating synthetic data and is thus called the *original dictionary*. By combining a small number of atoms of the original dictionary, a training dataset ($\mathbf{Y}_{\text{straight}}$) is produced. This training dataset is used to extract a *learned dictionary* with at least m atoms. The quality of the DLA is then evaluated by comparing how many atoms from the original dictionary have been recovered in the learned dictionary.

Here, the original dictionary \mathbf{U} of $m = 135$ normalized multivariate atoms is created from white uniform noise and the length of the atoms is $n = 20$ samples. The training set $\mathbf{Y}_{\text{straight}}$ is composed of $q = 2000$ signals of length n and it is synthetically generated from dictionary \mathbf{U} . Each training signal is generated as of the sum of three atoms, the coefficients and the atom indices being randomly drawn. At first, no noise is added to the generated signals for the set of experiments conducted in Section 7.3. An analysis of the noise sensitivity is proposed in Section 7.4. The dictionary initialization is made on the training set, and the learned dictionary $\hat{\mathbf{U}}$ is returned after 80 iterations over the training set.

The experimental protocol was slightly changed to give the training set $\mathbf{Y}_{\text{rotation}}$. In this case, a rotation is applied for each atoms composing the training signals. The coefficient of the rotation matrix, see Equation (21), are drawn randomly.

7.2 Detection Rates and Metrics Calculus

In the experiments, a learned atom \hat{U}_i is considered as detected, or recovered, if its correlation value ν_i with its corresponding original atom U_i is greater than a chosen threshold s . In the multivariate case of M-DLA described in Equation (17), it is expressed as:

$$\nu_i = \left| \left\langle U_i, \hat{U}_i \right\rangle \right| \geq s . \quad (22)$$

where s is a threshold fixed at 0.99 or 0.97. The *detection rate* is defined as the percentage of the original dictionary atoms which are recovered in the learned dictionary and it is denoted as $t_{0.99}$ and $t_{0.97}$. This threshold-based approach is very common in the DLA community to obtain a quantitative measure on dictionaries, but it is not a metric. Remark that this detection

¹In the context of rotation invariance, this experiment can be viewed as the extension of Barthélemy et al. (2013b) in dimension $\rho = 10$.

rate is invariant to sign: \hat{U}_i is considered recovered if it is close to the original atom U_i or to its opposite $-U_i$. The invariance to scale is already obtained since atoms are normalized.

This is naturally extended to include rotation invariance, as shown in Equation (21) for nDRI-DLA setup, by computing the correlation ν_i between U_i and \hat{U}_i after nD registration:

$$\nu_i \geq s \quad \text{with} \quad (\nu_i, \cdot) = \text{nD_Registration}(\hat{U}_i, U_i) . \quad (23)$$

Due to the nD registration step, this detection rate is invariant to rotation: \hat{U}_i is considered recovered if it is a rotated version of the original atom U_i .

In the following, two multivariate DLA algorithms are tested: M-DLA and the rotation invariant nDRI-DLA. The M-DLA is tested on the dataset $\mathbf{Y}_{\text{straight}}$, which does not include rotation of original atoms. The nDRI-DLA is evaluated on the two datasets $\mathbf{Y}_{\text{straight}}$ and $\mathbf{Y}_{\text{rotation}}$. The assessment of nDRI-DLA for both datasets relies on the detection rate based on Equation (23). The evaluation of the M-DLA detection rate is based on Equation (22).

Thanks to the Equation (16), it is possible to define real set-metric between dictionaries \mathbf{U} and $\hat{\mathbf{U}}$. With set-metrics, one important choice is the ground metric. For the multivariate dictionaries presented here, any distance described in Section 2.3 could be selected as ground distance for Hausdorff or Wasserstein metric. The ground metric allows to compute the distance between a learned atom and its associated atom in the original dictionary. For now on, the notation is changed for the sake of clarity to indicate the set-metric as a subscript and the ground distance as a superscript, for example d_{W}^{c} is the Wasserstein distance based on chordal ground distance.

To emphasize the importance of the ground distance choice and to clarify the results, only two ground distances are considered in this section. The first ground metric is the chordal distance d^{c} , defined in Equation (2), chosen for its invariance to linear combination, thus including rotations. As all principal angle-based ground metric share these properties, the chordal have been selected only because it is a widely accepted choice for Grassmannian manifold.

The second ground metric is a simple Frobenius-based distance d^{f} , expressed for multivariate atoms U_i and \hat{U}_j as:

$$\left(d^{\text{f}}(U_i, \hat{U}_j)\right)^2 = \|\hat{U}_j - U_i\|_F^2 = 2 - 2 \text{trace}(\hat{U}_j^T U_i) = 2 \left(1 - \langle U_i, \hat{U}_j \rangle\right) ,$$

under the classical assumption that the dictionary atoms have unit norm in the Frobenius sense, $\|U_i\|_F = \|\hat{U}_j\|_F = 1$. This distance is a metric in Euclidean space but is not in Grassmannian spaces, it is thus not described in Section 2.3. The distance d^{f} is related to the detection rate of Equation (22),




Figure 3: Evolution of detection rates and dictionary distances for M-DLA applied on $\mathbf{Y}_{\text{straight}}$ as a function of the learning iterations.

but without the sign invariance: \hat{U}_j is not considered recovered if it is close to $-U_i$. This distance is not invariant to linear transforms and hence to rotation and sign change.

The distances d^c and d^f served as ground distance for Hausdorff and Wasserstein set-metrics. For the latter, it is parametrized with $p = 1$, see Equation (12), and also known as Earth Mover’s distance or Kantorovich distance. Also, the measures are uniform on the whole support. We investigate here the four possible combinations: Hausdorff over chordal (d_H^c) or over Frobenius-based distance (d_H^f) and Wasserstein over chordal (d_W^c) or Frobenius (d_W^f). The set-metric distance are rescaled to be visually comparable to detection rates which are given in percentage:

$$\tilde{d}^c(\mathbf{U}, \hat{\mathbf{U}}) = \frac{\sqrt{\rho} - d^c(\mathbf{U}, \hat{\mathbf{U}})}{\sqrt{\rho}} \times 100,$$

$$\tilde{d}^f(\mathbf{U}, \hat{\mathbf{U}}) = \frac{\sqrt{2} - d^f(\mathbf{U}, \hat{\mathbf{U}})}{\sqrt{2}} \times 100.$$

7.3 Dictionary Recovering

The M-DLA is evaluated on a dataset generated from a combination of the original atoms, that is dataset $\mathbf{Y}_{\text{straight}}$ which does not include rotations. The results are presented on Figure 3, where the recovering rate of the learned dictionary, in term of distances or detection rates, is assessed for each iteration.

The detection rates $t_{0.99}$ or $t_{0.97}$ raises abruptly from 0 to around 70% between iterations 5 and 20 and shows non-smooth variations around stable




Figure 4: Same as Figure 3 with nDRI-DLA on $\mathbf{Y}_{\text{straight}}$: it is a multivariate dictionary learning with rotation invariance (nDRI-DLA) applied on a dataset without rotation of the signals ($\mathbf{Y}_{\text{straight}}$).

value, 67% for $t_{0.99}$ and 72% for $t_{0.97}$, for the rest of the iterations. The Wasserstein set-metrics offer a better picture, this distance is more smooth and more stable than the detection rates. This metric captures the fact that as the learned dictionary atoms are initialized with training signals of the dataset, the distance between the learned and the original dictionaries is not null, which is not the case for the detection rates. When using the Wasserstein set-metrics, the choice of ground distance has only a limited influence: the Frobenius-based metric stabilizes earlier than the chordal-based one. For the Hausdorff set-metrics, the experimental results confirm a known fact: the Hausdorff distances are less sensitive to internal variations inside the considered sets as they capture the distance of the most extreme point. Frobenius-based Hausdorff metric does not capture the changes caused by the learning algorithm, it stays with values around 40% for all iterations, whereas the chordal-based one shows a similar pattern to the Wasserstein distances but with a smaller amplitude.

The results obtained with nDRI-DLA on dataset $\mathbf{Y}_{\text{straight}}$ are shown on Figure 4. The set-metric – Wasserstein on chordal and Frobenius or Hausdorff on chordal – indicate that the algorithm slowly improves during the first 20 iterations before stabilizing. As observed on M-DLA, the Hausdorff distance based on chordal shows no sensitivity to the dictionary evolution during the learning phase. Another remark concerns the ground distance, here Frobenius or chordal, which has a major influence on the set-metric distance: the Wasserstein and Hausdorff distance based on chordal give very similar results. The detection rates are null for the whole learning phase and




Figure 5: Same as Figure 4 with nDRI-DLA on $\mathbf{Y}_{\text{rotation}}$: it is a multivariate dictionary learning with rotation invariance (nDRI-DLA) applied on a dataset with rotation of the signals ($\mathbf{Y}_{\text{rotation}}$).

thus completely uninformative.

The poor results of nDRI-DLA on $\mathbf{Y}_{\text{straight}}$ shown on Figure 4 are caused by the small size of the dataset given the large dimension of the multivariate components to learned: only $q = 2000$ signals for learning $m = 135$ atoms with $\rho = 10$ components. This is enough to achieve decent performances with M-DLA, but as nDRI-DLA should learn the rotation matrix on top of recovering the atoms and their coefficient, nDRI-DLA failed to obtain correct performance on dataset $\mathbf{Y}_{\text{straight}}$. However, with dataset $\mathbf{Y}_{\text{rotation}}$ each input signal contains three atoms from the original dictionary with different rotation coefficient, it is thus easier for nDRI-DLA to recover the atoms and learn the rotation matrices, as it is demonstrated on Figure 5.

With dataset $\mathbf{Y}_{\text{rotation}}$, nDRI-DLA performances are shown on Figure 5. The detection rate with $t_{0.99}$ stays at 0, whereas $t_{0.97}$ abruptly increase from iteration 25 and displays large non-smooth fluctuations due to its binary nature. Moreover, for $t_{0.99}$ very few atoms are considered as recovered, whereas for $t_{0.97}$ (which remains a selective rate) around 75% of atoms are considered as recovered. This highlights the extreme sensitivity of the detection rate with respect to the threshold value s . The Frobenius distance is affected by rotations and thus not suited for this setup: neither Wasserstein nor Hausdorff set-metric based on Frobenius distance are able to capture the evolution of the learned dictionary atoms. As explained previously the Frobenius-Wasserstein distance is not sensitive enough for the present setup. The chordal-based Wasserstein gives useful information on the convergence of the algorithm, which is smoothly increasing at each iteration. This example shows the importance of selecting an appropriate ground distance for set-metric, for example a rotation invariant distance such as the chordal is




Figure 6: Noise sensitivity analysis: averaged results for detection rates and distances with different noise level, that is 10, 20, 30 dB and no noise. (Top) M-DLA on dataset $\mathbf{Y}_{\text{straight}}$, (Center) nDRI-DLA on $\mathbf{Y}_{\text{straight}}$ (Bottom) nDRI-DLA on $\mathbf{Y}_{\text{rotation}}$.

required to correctly assessed the nDRI-DLA with dataset $\mathbf{Y}_{\text{rotation}}$.

7.4 Noise Sensitivity Analysis

For the next experiments, learning are repeated 10 times and Figure 6 shows the averaged detection rates and distances computed after 80 iterations. The previous experiments are replicated for different noise levels in the three setups: M-DLA on $\mathbf{Y}_{\text{straight}}$, nDRI-DLA on $\mathbf{Y}_{\text{straight}}$ and nDRI-DLA on

$\mathbf{Y}_{\text{rotation}}$. White Gaussian noise is added at several levels to obtain datasets with signal-to-noise ratio of 10, 20 and 30 dB, and a dataset is kept without adding noise. On each plot, to ease the comparison, the last column indicates the results obtained in the noiseless setting after 80 iterations, as in Figures 3, 4 and 5. The results are also averaged over 10 repetitions for the noiseless case.

In the M-DLA case, on the left of Figure 6, the detection rates and the distance are not affected by noise, showing the robustness of the learning algorithm on multivariate signals. These conclusions could be extended to the nDRI-DLA case on dataset $\mathbf{Y}_{\text{straight}}$, at the center of Figure 6. The original dictionary is poorly recovered, less than half of the atoms are identified, but the results are not affected by the presence of noise, even at high level. The detection rates $t_{0.97}$ and $t_{0.99}$ stay at 0 for all experiments and are thus completely inadequate for these evaluations.

Applying noise on the examples of dataset $\mathbf{Y}_{\text{rotation}}$ as the same effect on nDRI-DLA than applying random rotations on each original atoms: the atoms “pop out” from the noise, that is the atoms are the only stable signals existing in the dataset, all other parameters vary. It helps the algorithm to correctly reconstruct the original atoms, find their coefficient and their rotation matrix. This explains that the dictionary reconstructing in the 20 and 30 dB cases have results as good as the noiseless case. With only 10 dB, the noise yields small perturbations, corrupting the learned atoms and resulting in a decrease of recognition rate.

These experiments on synthetic dataset are aimed to demonstrate that direct applications of the previously described metrics offer immediate and handy tools for assessing dictionary learning algorithms. To exploit the full potential of the proposed metrics, we have chosen to propose experiments on multivariate datasets, but all these results could be applied directly on univariate datasets.

8 Experiments on EEG Signals

This section demonstrates some possible applications implied by the definition of metrics between dictionaries. We choose to present experiments on Brain-Computer Interface (BCI) because the spatio-temporal structure of the data offer, to some extent, an intuitive example of the application of multivariate set-metrics. Most of the BCI systems rely on electroencephalographic (EEG) measurements, which are electric variations caused by brain waves are recorded on the surface of the head with several electrodes. A dataset is thus an ensemble of timeseries, each of them representing the electric activity measured at a precise spatial location. Due to signal propagation and to physical properties of the head, two neighboring spatial locations have highly correlated signals, as it can be seen on Figure 7.




Figure 7: Example of an EEG during a 2 seconds motor imagery task. The electrical signals recorded from the 22 electrodes are plotted with a slight shift on the ordinate axis. Data from BCI Competition IV, set 2a, subject 7.

8.1 Experimental Protocol

Dictionary learning algorithms handling multivariate signals are able to capture the temporal as well as the spatial coherence of the EEG datasets (Barthélemy et al., 2013a). In this section, multivariate dictionaries are learned with M-DLA using the datasets of BCI Competition IV set 2a (Brunner et al., 2008; Tangermann et al., 2012). This is a motor imagery task, where 9 subjects have been instructed to imagine four classes of movements (left hand, right hand, tongue or feet) during two sessions conducted on different days. Each session consists of 288 trials (72 trials for each class) and a trial is the 3 seconds recording of $\rho = 22$ electrodes sampled at 250 Hz. For each session, for each subject and for each task, the M-DLA learns a dictionary of five shiftable patterns (giving a five time redundant dictionary), with a sparsity parameter $K = 1$. More details about this learning can be found in Barthélemy et al. (2013a).

In the following, two experiments are presented which rely on the set-metrics between dictionaries to investigate the properties of the BCI dataset thanks to clustering techniques. A part of the BCI community is organized and structured around the BCI Competition datasets, even if the individual variability of the subjects, both in term of inter-session evolution and in subject-specific characterization, is still largely unharnessed: between 20-30% of the BCI subjects obtain catastrophic results with the state of the art algorithms, a phenomenon called “BCI illiteracy” or “BCI-inefficiency” (Vidaurre and Blankertz, 2010; Hammer et al., 2012). Is it possible to detect subjects who are “BCI-inefficient”? Do the subjects in the different competition datasets have the same variability? How important is the inter-session variability, knowing that a separate session is often used as an evaluation set

for competition? All these questions require an objective measurement and set-metrics applied on learned dictionaries offer such a solution. The present experiments are intended to demonstrate that the proposed set-metric are ready to be applied and that their immediate application offer some interesting insight on BCI dataset.

The first experiment proposes to investigate the link between the clustering of subjects and their BCI performance, in order to gain a better understanding of the “BCI-inefficiency” phenomenon. The second experiment is oriented to investigate the subject-specific intersession variability and its influence on the dictionary learning.

8.2 Relation between Subject Clusters and BCI-Inefficiency

For these experiments, a dictionary is learned with M-DLA for each class of each subject: each subject is associated with four dictionaries (left hand, right hand, tongue and feet) learned over the first session of the dataset (used as a training set for the BCI Competition).

A distance matrix G^c between subject is computed for each class c : for a given class all the distances between subject’s dictionary are evaluated with the Wasserstein set-metric based on chordal distance. All the distances g_{ij}^c between subjects i and j are converted in Gaussian similarities s_{ij}^c with the following relation : $s_{ij}^c = \exp \left(-(g_{ij}^c)^2 / 2\sigma^2 \right)$ with $\sigma = 1$. For each class, clusters of subjects are gathered using affinity propagation (Frey and Dueck, 2007). Affinity propagation find the optimal number of clusters relying only on the similarities between subjects and the preference value of each subject, which the predisposition of each subject to become the exemplar of a cluster. Here, we apply a common approach where all subjects have an identical preference value equal to the median of G^c .

The obtained cluster for each class are then combined using the cluster ensembles approach defined by Strehl and Ghosh (2003). The results is a unique and global set of C clusters, C being a parameter, obtained by maximizing the shared mutual information of all classes. A thorough analysis shows that subjects are aggregated in 3 or 4 stable cluster ensembles, represented with different colors on the left side of Figure 8. On this Figure, the performance of each subject obtained with the state of the art algorithm (Ang et al., 2012) is shown on the ordinate-axis, the subjects are sorted according to their performance on the abscissa-axis. Using $C = 3$ cluster ensembles, it appears that the subject with the highest performance is always alone in a cluster, as the subject with the worst performance. All the other subjects are gathered in the same cluster. With $C = 4, 5$ or 6 , only the four cluster ensembles shown on Figure 8(Left) are found.

To have a better visualization of the relation between cluster, a hierarchical clustering is shown on the right side of Figure 8. It should be noticed




Figure 8: Left: Performance of the subjects from the BCI Competition IV-2a with state of the art algorithm from [Ang et al. \(2012\)](#), color indicates cluster ensemble obtained with consensus clustering on learned dictionaries. Right: Hierarchical clustering on the same dataset, using Hausdorff distance to determine linkage.

that instead of applying a common linkage criteria, such as the single linkage or the complete linkage, we rely here on the Hausdorff distance to evaluate the distance between sets of subjects. The results are very similar to those obtained with complete linkage. On the resulting dendrogram, it is visible that the subjects #2, #3, #9, #8, #5 and #4 belong to the same cluster and that each of the subjects #1, #7 and #6 constitute three separate clusters.

These experiments explicit new techniques to investigate the BCI inefficiency thanks to the set-metrics and the dictionary learning. In the BCI Competition IV dataset 2a it appears that most of the subjects seem to share a common profile except two extreme cases, a BCI inefficient and a BCI efficient ones. The possibilities offered by multivariate dictionary learning in conjunction with set-metric bring new opportunities to qualitatively assess datasets used in competitions and challenges. These approaches could help the community to propose more consistent and more complete benchmarks or evaluations.

8.3 Temporal Variability of Subject Brain Waves

In this section, the experiments aimed at investigating the temporal variability of a BCI subject, in other words how different are our brain waves from day to day? Appreciating this variability is crucial as it is very common to train a classifier with session “A” and to evaluate the same classifier on data acquired from session “B”, recorded several days after session “A”. This was the case for the BCI Competition IV dataset 2a. Intersession variability is also at the heart of the calibration problem: a common practice




Figure 9: Consensus clustering on the learned dictionaries with different ground metrics. The proportion of clusters specific to one session is plotted as a function of the number of clusters.

is to calibrate the electronic BCI system, including the classifiers, each day. Nonetheless a legitimate question when using a BCI system all day long is how often does it need to be calibrated? The intrinsic temporal evolution of brain waves is a long studied problem and yet no reference methodology has been found. The proposed experiments could help to properly design such methodology.

Here, a dictionary is learned with M-DLA for each subject, each class and each session. Thus for each session, “A” and “B”, there is 36 dictionaries learned, four by subject. The parameters used for M-DLA are the same than the previous section. The distances between dictionaries are computed with Wasserstein set-metric based on different ground distances: geodesic, chordal, Fubini-Study and Binet-Cauchy, defined respectively in Equations (1), (2), (3) and (5). As explained before, the clustering on dictionaries of each subject for a given class is obtained with the affinity propagation. The clusters for each class are then combined in cluster ensembles with the consensus algorithm of [Strehl and Ghosh \(2003\)](#). There are cluster ensembles for each ground distance and we also compute a global cluster ensembles gathering all classes and all ground distances for a given subject.

The 9 subjects are separated based on the time of their recordings, either in session “A” and “B”, resulting in 18 recordings to cluster. A cluster could gather recordings from sessions “A” or “B” and is said to belong to the session which has the most important representation among the recordings. For example, if a cluster gathers 3 recordings, two from session “A” and one from “B”, the cluster is said to belong to “A”. We propose to use as




Figure 10: The 9 subjects recorded during 2 different sessions (A and B). The spatial positions chosen according to Laplacian eigenmaps. The color indicates the 4 cluster ensembles obtained with consensus clustering on learned dictionaries. The thickness of the links is proportional to the distance between learned dictionaries.

measure of dissimilarity the proportion of recordings in the same cluster. A score of 1 indicates that all clusters contain either recordings from session “A” or, in the exclusive sense, from session “B”. The minimal score is 0 and indicates that each cluster contains as much recordings from session “A” and “B”. The chance level is at 0.5.

The Figure 9 shows the proportion of cluster’s recordings in the same session when the number C of cluster ensembles increases. The different metrics yield the similar results and the global cluster ensembles, computed from all ground distances, displays a high level of dissimilarity between the two sessions. Those results are stable with respect to the number of cluster ensembles C .

The Figure 10 proposes a different visualization of the same phenomenon. The 9 recordings of session “A” (one per subject) and the 9 ones of session “B” are represented with spatial positions drawn according to Laplacian eigenmaps (Belkin and Niyogi, 2003). By applying an eigenvalue decompo-

sition on the adjacency graph build from the distance matrix, the Laplacian eigenmaps embed the recording in \mathbb{R}^2 and thus offer an isometric representation. A thorough experimental analysis with the different metrics and the embedding produced by Laplacian eigenmaps leads us to choose $N = 10$ as the number of nearest neighbors picked by the Laplacian eigenmaps and discrete weights ($t = \infty$). Those parameters, used on the Fubini-Study distance matrix in the Figure 10, have been chosen for their illustrative potential and different parameters yield qualitatively similar results. The color code indicates how the recordings are separated into 4 cluster ensembles. The size of each link between recordings is proportional to the Wasserstein distance.

This figure shows that only two recordings (subjects #6 and #7 from session “A”) are associated with cluster ensembles containing mainly recordings from session “B”. On the one hand, it appears clearly that all the recordings from session “B” are clustered and are separated by small distances, with the notable exception of the subject #7. On the other hand, recordings from session “A” are all very distant from recording of session “B” and belong all to the same cluster, except for subjects #6 and #7 which are part of clusters of session “B”. The subject #7 is the one with the highest recognition rate, as shown on Figure 8: a tentative explanation is that the recordings of subject #7 are very similar between sessions “A” and “B” whereas all others subjects recording differ from one session to the other.

As a corollary, when using a measure reflecting the subject-dependence of the results, no clear results appear and the measure stays at the chance level. The measure used evaluates the dissimilarity between the subjects: each time the two sessions of a given subject are not in the same cluster, the indicator is decreased.

This demonstrates that there is a clear variation between sessions and that the current dictionary learning algorithm employed in these experiments, M-DLA, is sensitive to these session-dependent variations. The variations could be caused by the electronic system or could originate from the subject. A first conclusion is that a BCI system should be recalibrated at the beginning of each session to avoid a performance drop. The second conclusion concerns the dictionary learning algorithm: a training dataset including trials from several sessions could reinforce the robustness of the dictionary.

The last experiment concerns the classes of the motor imagery task. This time the distances are computed between classes for each subject, instead of computing distances between subjects for each class as in the previous experiment. The M-DLA is applied with the same parameters, the affinity propagation is applied on the classes for each subject and each ground metric (Binet-Cauchy, Fubini-Study, chordal and geodesic). Then the clusters of all subjects and metrics are merged in cluster ensembles with the consensus algorithm. The resulting cluster ensembles are shown on Figure 11: the hands and tongue imagery are gathered in the same cluster. The feet are set




Figure 11: Consensus clustering on motor imagery classes: the cluster ensemble reflects the neurobiological organization of the brain.

in a separate cluster. This clustering reflect the neurobiological organization of the brain: on the one hand the localization of the cortical areas devoted to hand and tongue movements are situated on the sides of the head, in the parietal areas. On the other hand, the surface of the motor cortex dedicated to the feet movements is situated on the interior side of the interhemispheric fissure, generating EEG components centered on the top of the head.

These experiments demonstrate that, properly equipped with a metric, the dictionary learning offer an efficient and robust approach to assess the intrinsic properties of multivariate datasets. The formalization of a metric space applicable to frames open several direct applications, as shown in this section.

9 Conclusion

This contribution rely on advances from algebraic geometry and frame theory to define suited metrics for multivariate dictionaries. It is the first contribution to define such metric spaces for learned dictionaries. The distance between dictionaries is computed as the distance between their subspaces, yielding a pseudo-metric space which is invariant to linear transformations, an essential property to compare multivariate dictionaries.

The interest of the described set-metric is shown through its direct application on two examples: a synthetic dataset and real dataset of EEG for Brain Computer Interfaces. The defined metrics allow to estimate empirically the convergence of a dictionary learning algorithm with a precision outperforming the classical measurements based on detection rates. Furthermore, they also allow to qualitatively assess a specific dataset as it is shown with the BCI experiments.

In experiments on synthetic and real datasets, it is shown that the set-metrics could be adapted to the specific properties of the considered data, as adding invariance to rotations by choosing a suited ground distance. Indeed, the choice of the ground distance offers a large amount of possibilities to

capture some specific properties linked to the properties of the underlying data. Ground metrics such as the chordal or the geodesic distance are a good default choice.

The chosen examples amount from multivariate data analysis, illustrating that the proposed metrics could be applied straightly on these complex problems. Nonetheless, all the work described in this contribution is also valid for the univariate case and operates with every dictionary learning algorithms. But in numerous applications, multivariate datasets are transformed to be processed by univariate algorithms, causing an important information loss. Metrics for multivariate dictionaries bring a framework to develop new and improved multivariate dictionary learning algorithms.

At least, an original connection is made between Grassmannian frames, compressed sensing and dictionary learning. This disjoint research topics try either to cover the entire space under consideration, as in packing problems, or to minimize the reconstruction error on the signal energy, as in dictionary learning, or to add coherence constraints on dictionary learning algorithms, as in hybrid approaches.

Acknowledgments

Sylvain Chevallier was partially supported by a grant from EADS Foundation, project Cerebraptic.

Appendix A.

In this appendix, we explicit the definitions of manifolds, matrix manifolds and Grassmannian Manifold.

Manifolds

Definition 27 *Let us consider an n -dimensional vector space \mathcal{V} with inner product $\langle \cdot, \cdot \rangle$ defined over the field \mathbb{R} . A manifold is a topological space \mathcal{M} of class C^k (differentiable up to k^{th} degree of differentiation) together with an atlas $\mathcal{A} = (\phi_i, Z_i)_{i \in I}$ on \mathcal{V} , with $I \subset \mathbb{N}$ an indexing set. An atlas is a collection of charts, a chart (ϕ_i, Z_i) being an homeomorphism of the open subset $Z_i \subset \mathcal{M}$ onto an open subset $\phi(Z_i) \subset \mathcal{V}$, that is (ϕ_i, Z_i) is a bijection, is continuous and its inverse ϕ^{-1} is continuous.*

A manifold is differentiable if it endows a globally defined differentiable structure. To induce a globally differentiable structure from the charts, the composition of intersecting charts must be a differentiable function and the charts are said C^k -compatible. For example if two charts overlap, the coordinates defined by a chart are required to be differentiable with respect to the coordinates defined in the other chart. Formally, any two charts $(\phi_i, Z_i), (\phi_j, Z_j)$ are C^k -compatible if

$$\phi_{ij} := \{\phi_i \circ \phi_j^{-1} |_{\phi(Z_i \cap Z_j)} : \phi_j(Z_i \cap Z_j) \rightarrow \phi_i(Z_i \cap Z_j)\} ,$$

and its inverse ϕ_{ji} are of class C^k . If \mathcal{M} is C^k for all k , \mathcal{M} is said smooth or C^∞ -manifold.

Matrix Manifolds

First, remark that if we let $\{e_i\}_{i=1}^n$ be a basis of \mathcal{V} , then the function

$$\phi : \mathcal{V} \rightarrow \mathbb{R}^n, u \mapsto [u_1, \dots, u_n]$$

such that $u = \sum_{i=1}^n u_i e_i$ is a chart of the set \mathcal{V} . As all the charts constructed in this way are compatible (or “smooth”), they form an atlas of the set \mathcal{V} , endowing \mathcal{V} with a manifold structure. Thus, every vector space is a linear manifold.

The set $\mathbb{R}^{n \times \varrho}$ with the usual sum and multiplication by a scalar is a vector space. Hence, it is endowed with a linear manifold structure and a chart of this manifold is

$$\phi : \mathbb{R}^{n \times \varrho} \rightarrow \mathbb{R}^{n\varrho}, U \mapsto \text{vec}(U) ,$$

where $\text{vec}(U)$ denotes the vectorization of the matrix U . The set $\mathbb{R}^{n \times \varrho}$ together with its linear manifold structure is denoted \mathbf{M} hereafter. The




Figure 12: Hierarchy of matrix manifolds

manifold \mathbf{M} is the set of all matrices in $\mathbb{R}^{n \times \varrho}$, that is the set of all linear applications $\mathcal{L} : \mathbb{R}^n \rightarrow \mathbb{R}^\varrho$ or equivalently the set $\mathbb{R}^\varrho \otimes \mathbb{R}^n$ defined as the tensor product of \mathbb{R}^ϱ and \mathbb{R}^n . Also, \mathbf{M} can be seen as an Euclidean space with the inner product $\langle U, W \rangle := \text{vec}(U)^T \text{vec}(W) = \text{trace}(U^T W)$ and the induced norm $\|U\|_F$.

Definition 28 Let n and ϱ a positive integer with $\varrho \leq n$. The nonorthogonal Stiefel manifold $\mathbf{St}(\varrho, n)$ is the set of all full rank matrices in \mathbf{M} , that is

$$\mathbf{St}(\varrho, n) := \{U \in \mathbf{M} : \text{rank}(U) = \varrho\} .$$

The column space \mathcal{U} of a matrix U in $\mathbf{St}(\varrho, n)$ defines the basis of a ϱ -dimensional subspace in \mathbb{R}^n , that is the ϱ column vectors (u_1, \dots, u_ϱ) of U span the subspace \mathcal{U} , thus $\text{span}(U) = \mathcal{U}$. For $n = \varrho$, the Stiefel manifold $\mathbf{St}(\varrho, \varrho)$ reduces to the general linear group $\mathbf{GL}(\varrho)$.

Stiefel manifold defined here is different from the orthogonal Stiefel manifold \mathbf{St}^* , which is the set of orthonormal n -by- ϱ matrices.

Definition 29 Let n and ϱ be two positive integers such that $\varrho \leq n$. The orthogonal Stiefel manifold $\mathbf{St}^*(\varrho, n)$ is defined as

$$\mathbf{St}^*(\varrho, n) := \{U \in \mathbf{M} : U^T U = I_\varrho\} ,$$

with I_ϱ the ϱ -by- ϱ identity matrix.

For $\varrho = 1$, $\mathbf{St}^*(1, n)$ reduces to the unit sphere \mathbf{S}^{n-1} and for $\varrho = n$, it reduces to the set of orthogonal matrices $\mathbf{O}(n)$, as shown on Figure 12.




Figure 13: Construction of the Grassmannian manifold. Here, the projection π verifies $\pi^{-1}(\pi(U)) = \{U\mathbf{GL}(\rho) : U \in \mathbf{GL}(\rho)\}$ and \tilde{u} is the projection of u on $\mathbf{Gr}(\rho, n)$ such that $u = \tilde{u} \circ \pi$.

Grassmannian Manifold

The Grassmannian manifold $\mathbf{Gr}(\rho, n)$ is a quotient manifold, such as an element \mathcal{U} of $\mathbf{Gr}(\rho, n)$ can also be characterized as a ρ -dimensional subspace in \mathbb{R}^n . A matrix U is said to span \mathcal{U} if $\text{span}(U) = \mathcal{U}$, and the span of U is an element of $\mathbf{Gr}(\rho, n)$ only if $U \in \mathbf{St}(\rho, n)$. There is an infinite number of matrices in $\mathbf{St}(\rho, n)$ which span an element \mathcal{U} of $\mathbf{Gr}(\rho, n)$, as illustrated on Figure 13.

Definition 30 Given a matrix $U \in \mathbf{St}(\rho, n)$, the set of the matrices that have the same span as U is defined as

$$U\mathbf{GL}(\rho) := \{UA : A \in \mathbf{GL}(\rho)\} ,$$

where $\mathbf{GL}(\rho)$ is the set of ρ -by- ρ invertible matrices. The Grassmannian manifold $\mathbf{Gr}(\rho, n)$ is defined as the orbit space of

$$\mathbf{St}(\rho, n)/\mathbf{GL}(\rho) := \{U\mathbf{GL}(\rho) : U \in \mathbf{St}(\rho, n)\} ,$$

with the group action $\mathbf{St}(\rho, n) \times \mathbf{GL}(\rho) \rightarrow \mathbf{St}(\rho, n), (U, A) \mapsto UA$.

In the case $n = 2$ and $\rho = 1$, an element of a Grassmannian manifold defines an equivalence line in the plane, that is a set of points in $\mathbf{St}(\rho, n)$ which are invariant under the group $\mathbf{GL}(\rho)$. In a more general case, each element \mathcal{U} of a $\mathbf{Gr}(\rho, n)$ defines a subspace and for each \mathcal{U} corresponds an equivalent class of matrices that span \mathcal{U} .

References

- M. Aharon. *Overcomplete Dictionaries for Sparse Representation of Signals*. PhD thesis, Technion - Israel Institute of Technology, 2006.
- M. Aharon, M. Elad, and A.M. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Processing*, 54:4311–4322, 2006.

- A. Aldroubi. Portraits of Frames. *Proceedings of the American Mathematical Society*, 123(6):1661–1668, 1995.
- K.K. Ang, Z.Y. Chin, C. Wang, C. Guan, and H. Zhang. Filter bank common spatial pattern algorithm on BCI competition IV datasets 2a and 2b. *Frontiers in Neuroscience*, 6:1–9, 2012.
- L. Applebaum, S.D. Howard, S. Searle, and R. Calderbank. Chirp sensing codes: Deterministic compressed sensing measurements for fast recovery. *Applied and Computational Harmonic Analysis*, 26(2):283–290, 2009.
- R. Balan. Equivalence Relations and Distances between Hilbert Frames. *Proceedings of the American Mathematical Society*, 127(8):2353–2366, 1999.
- Q. Barthélemy, A. Larue, A. Mayoue, D. Mercier, and J.I. Mars. Shift & 2D rotation invariant sparse coding for multivariate signals. *IEEE Trans. Signal Processing*, 60:1597–1611, 2012.
- Q. Barthélemy, C. Gouy-Pailler, Y. Isaac, A. Souloumiac, A. Larue, and J.I. Mars. Multivariate temporal dictionary learning for EEG. *Journal of Neuroscience Methods*, in press, 2013a.
- Q. Barthélemy, A. Larue, and J.I. Mars. Decomposition and dictionary learning for 3D trajectories. Technical report, CEA, 2013b.
- A. Beck and M. Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM J. Img. Sci.*, 2(1):183–202, 2009.
- M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.*, 15(6):1373–1396, 2003.
- C. Brunner, R. Leeb, G.R. Müller-Putz, A. Schlögl, and G. Pfurtscheller. BCI Competition 2008 - Graz data set A, 2008.
- D. Burago, Y. Burago, and S. Ivanov. *A course in metric geometry*. American Mathematical Society Providence, 2001.
- E.J. Candès and D.L. Donoho. New tight frames of curvelets and optimal representations of objects with piecewise C2 singularities. *Comm. Pure Appl. Math.*, 57(2):219–266, 2004.
- E.J. Candès, J.K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1207–1223, 2006.

- E.J. Candès, J.K. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Information Theory*, 52(2):489–509, 2006.
- P.G. Casazza and G. Kutyniok. *Frames of subspaces*, pages 87–114. Wavelets, Frames and Operator Theory. AMS, 2004.
- Y. Chikuse. *Statistics on Special Manifolds*. Springer, 2003.
- O. Christensen. *An introduction to frames and Riesz bases*. Birkhauser, 2003.
- J.H. Conway, R.H. Hardin, and N.J.A. Sloane. Packing lines, planes, etc.: Packings in grassmannian spaces. *Experimental Mathematics*, 5(2):139–159, January 1996.
- I. Daubechies, A. Grossmann, and Y. Meyer. Painless nonorthogonal expansions. *Journal of Mathematical Physics*, 27(5):1271–1283, 1986.
- G. Davis. *Adaptive Nonlinear Approximations*. PhD thesis, New York University, 1994.
- R.A. DeVore. Deterministic constructions of compressed sensing matrices. *Journal of Complexity*, 23(4-6):918–925, 2007.
- I.S. Dhillon, R.W. Heath Jr., T. Strohmer, and J.A. Tropp. Constructing Packings in Grassmannian Manifolds via Alternating Projection. *Experimental Mathematics*, 17(1):9–35, 2008.
- D.L. Donoho. Compressed sensing. *IEEE Trans. Information Theory*, 52:1289–1306, 2006.
- R.J. Duffin and A.C. Schaeffer. A class of nonharmonic Fourier series. *Trans. Amer. Math. Soc.*, 72(2):341–366, 1952.
- A. Edelman, T.A. Arias, and S.T. Smith. The Geometry of Algorithms with Orthogonality Constraints. *SIAM J. Matrix Anal. Appl.*, 20(2):303–353, 1999.
- B. Efron, T. Hastie, L. Johnstone, and R. Tibshirani. Least Angle Regression. *The Annals of Statistics*, 32(2):407–451, 2004.
- K. Engan, S.O. Aase, and J.H. Husøy. Multi-frame compression: theory and design. *Signal Process.*, 80:2121–2140, 2000.
- K. Engan, K. Skretting, and J.H. Husøy. Family of iterative LS-based dictionary learning algorithms, ILS-DLA, for sparse signal representation. *Digit. Signal Process.*, 17:32–49, 2007.

- B.J. Frey and D. Dueck. Clustering by passing messages between data points. *Science*, 315:972–976, 2007.
- G.H. Golub and C.F. van Loan. *Matrix Computations*. The Johns Hopkins University Press, 3rd edition, 1996.
- G.H. Golub and H. Zha. The canonical correlations of matrix pairs and their numerical computation. *Institute for Mathematics and Its Applications*, 69:27–49, 1995.
- R. Gribonval and K. Schnass. Dictionary Identification–Sparse Matrix-Factorization via ℓ_1 -Minimization. *IEEE Trans. Information Theory*, 56(7):3523–3539, 2010.
- R. Gribonval, H. Rauhut, K. Schnass, and P. Vandergheynst. Atoms of all channels, unite! average case analysis of multi-channel sparse recovery using greedy algorithms. Technical Report PI-1848, IRISA, 2007.
- R. Grosse, R. Raina, H. Kwong, and A.Y. Ng. Shift-Invariant Sparse Coding for Audio Classification. In *Proc. Conf. on Uncertainty in Artificial Intelligence UAI*, 2007.
- J. Hamm and D.D. Lee. Grassmann discriminant analysis: a unifying view on subspace-based learning. In *Proc. of ICML*, pages 376–383. ACM, 2008.
- E.M. Hammer, S. Halder, B. Blankertz, C. Sannelli, T. Dickhaus, S. Kleih, K.-R. Müller, and A. Kübler. Psychological predictors of SMR-BCI performance. *Biological Psychology*, 89(1):80–86, 2012.
- H. Hotelling. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377, 1936.
- R. Jenatton, R. Gribonval, and F. Bach. Local stability and robustness of sparse dictionary learning in the presence of noise, 2012.
- S. Kobayashi and K. Nomizu. *Foundations of Differential Geometry*. Wiley-Interscience, 1969.
- G. Kutyniok, A. Pezeshki, R. Calderbank, and T. Liu. Robust dimension reduction, fusion frames, and grassmannian packings. *Applied and Computational Harmonic Analysis*, 26(1):64–76, 2009.
- S. Lesage. *Apprentissage de dictionnaires structurés pour la modélisation parcimonieuse des signaux multicanaux*. Thèse de Doctorat, Université de Rennes, 2007.
- Y.M. Lui. Advances in matrix manifolds for computer vision. *Image and Vision Computing*, 30(6-7):380–388, 2012.

- B. Maill  , R. Gribonval, F. Bimbot, M. Lemay, P. Vandergheynst, and J.M. Vesin. Dictionary learning for the sparse modelling of atrial fibrillation in ECG signals. In *IEEE ICASSP*, pages 465–468, 2009.
- J. Mairal, M. Elad, and G. Sapiro. Sparse representation for color image restoration. *IEEE Trans. Image Processing*, 17:53–69, 2008.
- J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Non-local sparse models for image restoration. In *IEEE ICCV*, pages 2272–2279, 2009.
- J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online learning for matrix factorization and sparse coding. *J. Mach. Learn. Res.*, 11:19–60, 2010.
- S. Mallat. *A Wavelet Tour of Signal Processing, Second Edition (Wavelet Analysis & Its Applications)*. Academic Press, 1999.
- S.G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Processing*, 41:3397–3415, 1993.
- Y. Meyer. *Wavelets and Operators: Volume 1*. Cambridge University Press, 1995.
- J.W. Milnor and J.D. Stasheff. *Characteristic Classes.(AM-76)*, volume 76. Princeton University Press, 1974.
- D.G. Mixon, C.J. Quinn, N. Kiyavash, and M. Fickus. Equiangular tight frame fingerprinting codes. In *IEEE ICASSP*, pages 1856–1859, 2011.
- G. Monaci, P. Jost, P. Vandergheynst, B. Maill  , S. Lesage, and R. Gribonval. Learning multimodal dictionaries. *IEEE Trans. Image Processing*, 16:2272–2283, 2007.
- G. Monaci, P. Vandergheynst, and F.T. Sommer. Learning bimodal structure in audio-visual data. *IEEE Trans. Neural Networks*, 20:1898–1910, 2009.
- Y. Moudden, J. Bobin, J.-L. Starck, and J. Fadili. Dictionary learning with spatio-spectral sparsity constraints. In *SPARS '09*, 2009.
- Y.C. Pati, R. Rezaiifar, and P.S. Krishnaprasad. Orthogonal Matching Pursuit: recursive function approximation with applications to wavelet decomposition. In *Proc. Asilomar Conf. on Signals, Systems and Comput.*, 1993.
- A. Rakotomamonjy. Surveying and comparing simultaneous sparse approximation (or group-lasso) algorithms. *Signal Process.*, 91:1505–1526, 2011.
- K. Skretting and K. Engan. Recursive least squares dictionary learning algorithm. *IEEE Trans. Signal Processing*, 58:2121–2130, 2010.

- K. Skretting and K. Engan. Learned dictionaries for sparse image representation: properties and results. In *SPIE Conference*, volume 8138, 2011.
- N. Strawn. Optimization over finite frame varieties and structured dictionary design. *Applied and Computational Harmonic Analysis*, 32:413–434, 2012.
- A. Strehl and J. Ghosh. Cluster ensembles — a knowledge reuse framework for combining multiple partitions. *J. Mach. Learn. Res.*, 3:583–617, 2003.
- T. Strohmer and R.W. Heath Jr. Grassmannian frames with applications to coding and communication. *Applied and Computational Harmonic Analysis*, 14(3):257–275, 2003.
- P. Tammes. *On the origin of number and arrangement of the places of exit on the surface of pollen-grains*. PhD thesis, Faculty of Groningen, 1930.
- M. Tangermann, K.-R. Müller, A. Aertsen, N. Birbaumer, C. Braun, C. Brunner, R. Leeb, C. Mehring, K.J. Miller, G. Müller-Putz, G. Nolte, G. Pfurtscheller, H. Preissl, G. Schalk, A. Schlögl, C. Vidaurre, S. Waldert, and B. Blankertz. Review of the BCI Competition IV. *Frontiers in Neuroscience*, 6(55), 2012.
- R. Tibshirani. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
- I. Tošić and P. Frossard. Dictionary learning. *IEEE Signal Processing Magazine*, 28:27–38, 2011a.
- I. Tošić and P. Frossard. Dictionary learning for stereo image representation. *IEEE Trans. Image Processing*, 20:921–934, 2011b.
- J.A. Tropp. Greed is good: algorithmic results for sparse approximation. *IEEE Trans. Information Theory*, 50(10):2231–2242, 2004.
- J.A. Tropp and S.J. Wright. Computational methods for sparse solution of linear inverse problems. *Proceedings of the IEEE*, 98:948–958, 2010.
- D. Vainsencher, S. Mannor, and A.M. Bruckstein. The sample complexity of dictionary learning. *J. Mach. Learn. Res.*, 12:3259–3281, 2011.
- C. Vidaurre and B. Blankertz. Towards a cure for BCI illiteracy. *Brain Topography*, 23(2):194–198., 2010.
- S.V.N. Vishwanathan and A.J. Smola. Binet-Cauchy kernels. In *Advances in Neural Information Processing Systems (NIPS)*, 2004.
- L. Wang, X. Wang, and J. Feng. Subspace distance analysis with application to adaptive Bayesian algorithm for face recognition. *Pattern Recognition*, 39(3):456–464, 2006.

- L. Welch. Lower bounds on the maximum cross correlation of signals (corresp.). *IEEE Trans. Information Theory*, 20:397–399, 1974.
- L. Wolf and A. Shashua. Learning over sets using kernel principal angles. *J. Mach. Learn. Res.*, 4:913–931, 2003.
- Y.C. Wong. Differential geometry of Grassmann manifolds. *PNAS*, 57(3):589–594, 1967.
- M. Yaghoobi, L. Daudet, and M.E. Davies. Parametric dictionary design for sparse coding. *IEEE Trans. Signal Processing*, 57:4800–4810, 2009.